

10 Windows 8.1 Tips

# Windows IT Pro

A PENTON PUBLICATION

FEBRUARY 2014 | [WINDOWSITPRO.COM](http://WINDOWSITPRO.COM) | WE'RE IN IT WITH YOU

Implement

# QoS

in Windows Server 2012

Manage Active Directory  
with PowerShell

Data Protection Manager  
for Virtualized Workloads

6 Problems with  
Failover Clusters

Office 365  
Message Encryption



# Top 5 Tools that IT Pros Love... and Cats Too!

Get Them All  
[netwrix.com/5\\_tools](http://netwrix.com/5_tools)

## 1. Netwrix Auditor

Configuration auditing for IT systems

Netwrix Auditor monitors your environment 24x7 and helps you answer the critical questions of 'who changed what, when and where'. Designed to support the widest variety of IT systems and applications, Netwrix Auditor won't let any change slip away or stay unnoticed and drastically improves IT infrastructure visibility and internal security. Audited systems include Active Directory, Group Policy, Exchange, Windows Server, File Server, SQL Server, Event Logs & Syslog events, VMware, SharePoint and more!

## 2. Netwrix Password Manager

Self-service password reset

Password Manager allows users to reset forgotten passwords, troubleshoot account lockouts and unlock their accounts manually, through a convenient, web-based, self-service portal and integration with the standard Windows logon procedure.

## 3. Netwrix Account Lockout Examiner

Troubleshooting of account lockouts

Account Lockout Examiner automatically alerts the help-desk staff on lockout events and launches a troubleshooting process, scanning through system services, mapped network drivers, scheduled tasks and other places. The accounts can be unlocked via the unified console, web-based interface or even a mobile device.

## 4. Netwrix Disk Space Monitor

Automatic monitoring of low disk space

Disk Space Monitor alerts administrators in real-time when disk space falls below certain threshold on one or more network servers. This freeware tool centrally monitors multiple servers for free disk space and sends alerts by e-mail.

## 5. Netwrix Bulk Password Reset

One-click reset of local Administrator password on multiple computers

Bulk Password Reset comes handy for changing multiple local account passwords in bulk and across multiple computers. You can specify an entire domain, select a specific subset of computers, or tell the tool to use a computer list from a text file.



# A 360° DESKTOP AND MDM SOFTWARE DEVICE MANAGEMENT SIMPLIFIED



Patch Management . Software Deployment . Asset Management .  
Security Management . Profile Management . Application Management



# COVER STORY ▼

## Implementing QoS in Windows Server 2012 39

— John Savill

With Windows Server 2012, QoS introduces relative minimums that simplify the process and help you maximize resource utilization. Find out more about software, hardware, and virtualization QoS and how to manage them in your organization.

### Features

#### 51 Managing Active Directory with PowerShell

Darren Mar-Elia

#### 65 Data Protection Manager for Virtualized Workload Protection

John Savill

#### 75 Office 365 Message Encryption

Tony Redmond

### Products

#### 83 New & Improved

### Interact

#### 78 Ask the Experts

### In Every Issue

#### 87 Advertiser Directory

#### 87 Directory of Services

#### 87 Vendor Directory

### Chat with Us



Facebook



Twitter



LinkedIn



# Columns



8 [Need to Know](#)

## **Planning for a Post-PC World**

Paul Thurrott



14 [Windows Power Tools](#)

## **From One-Liner to ForEach One-Liner**

Mark Minasi



18 [Top 10](#)

## **Make Windows 8.1 Work Like Windows 7**

Michael Otey



22 [Enterprise Identity](#)

## **Windows Azure Active Directory Updates in 2013**

Sean Deuby



27 [What Would Microsoft Support Do?](#)

## **Six Common Problems with Failover Clusters**

John Marlin

## Editorial

Vice President, Content & User Engagement:  
Joe Panettieri

Editorial Director: Megan Keller

Editor-in-Chief: Amy Eisenberg

Senior Technical Director: Michael Otey

Technical Director: Sean Deuby

Senior Technical Analyst: Paul Thurrott

IT Community Manager: Rod Trent

Systems Management, Exchange Server &  
Outlook: Jason Bovberg

Scripting, Developer Content:

Blair Greenwood

SQL Server: Jayleen Heft

SharePoint, Active Directory, Security,

Virtualization: Caroline Marwitz

Managing Editor: Lavon Peters

## Senior Contributing Editors

David Chernicoff, Mark Minasi,  
Tony Redmond, Paul Robichaux,  
Mark Russinovich, John Savill

## Contributing Editors

Alex K. Angelopoulos, Michael Dragone,  
Jeff Felling, Brett Hill, Dan Holme,  
Darren Mar-Elia, Eric B. Rux,  
William Sheldon, Curt Spanburgh,  
Bill Stewart, Orin Thomas, Douglas Toombs,  
Ethan Wilansky

## Art & Production

Senior Graphic Designer: Matt Wiebe  
Group Production Manager:  
Julie Jantzer-Ward  
Project Manager: Adriane Wineinger  
Graphic Specialist: Karly Prickett

## Advertising Sales

Strategic Accounts Director:  
Chrissy Ferraro • 970-203-2883

Account Executives:

Megan Key • 970-203-2844

Barbara Ritter • 858-367-8058

Cass Schulz • 858-357-7649

## Client Services

Senior Client Services Manager:  
Michelle Andrews • 970-613-4964

## Marketing & Circulation

Customer Service • 800-793-5697

Vice President, User Marketing &

Marketing Analytics: Tricia Syed

Marketing Director: Amy Connell

## Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

## Corporate

Chief Executive Officer:

David Kieselstein

Chief Financial Officer/Executive Vice  
President: Nicola Allais



## List Rentals

Sarah Nowowiejski

## Reprints

Reprint Sales:

Wright's Media • 877-652-5295

*Windows IT Pro*, February 2014, Issue No. 234,  
ISSN 1552-3136. *Windows IT Pro* is published monthly by  
Penton. Copyright ©2014 Penton. All rights reserved. No  
part of this publication may be reproduced or distributed  
in any way without the written consent of Penton.

*Windows IT Pro*, 748 Whalers Way, Fort Collins, CO 80525,  
800-621-1544 or 970-663-4700. Customer Service:  
800-793-5697.

We welcome your comments and suggestions about the  
content of *Windows IT Pro*. We reserve the right to edit all  
submissions. Letters should include your name and  
address. Please direct all letters to [letters@windowsitpro.com](mailto:letters@windowsitpro.com). IT pros interested in writing for *Windows IT Pro* can  
submit articles at [windowsitpro.com/node/submission/](http://windowsitpro.com/node/submission/)  
article.

Program Code: Unless otherwise noted, all programming  
code in this issue is ©2014, Penton, all rights reserved.  
These programs may not be reproduced or distributed  
in any form without permission in writing from the  
publisher. It is the reader's responsibility to ensure  
procedures and techniques used from this publication are  
accurate and appropriate for the user's installation. No  
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®  
are trademarks or registered trademarks of Microsoft  
Corporation in the United States and/or other countries  
and are used by Penton, under license from owner.  
*Windows IT Pro* is an independent publication not  
affiliated with Microsoft Corporation. Microsoft  
Corporation is not responsible in any way for the editorial  
policy or other contents of the publication.

# WindowsITPro

# RSA<sup>®</sup> CONFERENCE 2014

FEBRUARY 24 – 28 | MOSCONE CENTER | SAN FRANCISCO

Share. Learn. Secure.

Capitalizing on Collective Intelligence



Register by February 21 for  
savings off the onsite price

2 Expos | 350+ Exhibitors | 21 Tracks  
300+ Sessions | 17 Keynotes



Closing Keynote Speaker

**STEPHEN COLBERT**

Award-winning host and  
executive producer of  
"The Colbert Report" and  
New York Times best selling author

Experience new ways of learning  
with these exciting opportunities:

- > **NEW** – The **Sandbox** featuring *Innovation Sandbox* and *The Most Innovative Company*
- > **Flash Talks** Powered by PechaKucha
- > **Association Events** and **Track Sessions**
- > **CE Credits**

FOLLOW US ON:

#RSAC



Register Now! [www.rsaconference.com/itpro](http://www.rsaconference.com/itpro)

Global Diamond Sponsors



Global Platinum Sponsors



Global Gold Sponsors



Platinum Sponsors



Gold Sponsors





# Here it is!



Mark Woffinden  
2011 CEO Financial Report  
Sent: Fri 12/30/2011 12:03 PM  
To: Board of Directors

**Even if your Exchange Server is like the lost sands of the Sahara,  
StorageCraft can help you recover each and every grain.**

Simplify search, recovery, and migration of Exchange servers. StorageCraft® Granular Recovery for Exchange and StorageCraft ShadowProtect® Granular Recovery for Exchange both support Microsoft® Exchange Server 2013 and allow you to recover or migrate individual or multiple mailboxes from Exchange Database (EDB) files quickly and easily.

**Use with any backup solution:**

**STORAGECRAFT®**  
GRANULAR RECOVERY FOR EXCHANGE

**Use with StorageCraft ShadowProtect:**

STORAGECRAFT.  
**SHADOWPROTECT®**  
GRANULAR RECOVERY FOR EXCHANGE

Whether used with a StorageCraft backup image or a backup from another provider,  
you'll keep your business running after any email disaster—great or small.



Download your  
**FREE 30-Day Trial**  
[www.StorageCraft.com/WINGRE](http://www.StorageCraft.com/WINGRE)



  
**STORAGECRAFT®**  
*Backup Fast, Recover Faster*

# Planning for a Post-PC World



## Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



As 2014 begins, Microsoft faces a more difficult competitive landscape than it has since Windows rose to prominence over 20 years ago. Increasingly irrelevant in personal computing, the firm must find its way in what many are calling the post-PC world.

What is Microsoft's role in this post-PC world? And how will the changes that are rocking personal computing affect those who have historically relied on Windows and Microsoft for their livelihoods?

That we're in a post-PC world is, of course, a matter of semantics. Microsoft has countered this claim, which was made famous by former Apple CEO Steve Jobs, by asserting that we're actually in a "PC-plus era," one in which PCs are used alongside other digital devices such as tablets and smartphones.

While arguments could be made for either term, the simple truth is this: Sometime soon, potentially this year, PCs will be the slowest-selling of the three mainstream computing device types. It's only a matter of time before tablet usage surpasses that of PCs as well. (Smartphones are already well ahead.)

The PC isn't out. But it's on the way down.

## Post-PC World: Smartphones, Tablets, and PCs

The decline of the PC is a problem because while Microsoft Windows has long dominated the PC market, it has found little foothold in tablets or smartphones. These markets are dominated by Android, which many have described as the Windows of the device world. In 2013, Android accounted for over 80 percent of all smartphones sold and about 65 percent of all tablets.

If you examine the sales of all of these devices together—a market we might describe as "personal computing devices"—Windows accounted for only about 14 percent of that market in 2013, compared to 38 percent for Android and 11.6 percent for Apple platforms

(Mac plus iOS). And while sales of Windows on non-PC devices are expected to rise over the next few years, that increase will occur much more slowly than will the rate for Android and iOS devices. Since PC sales are going down each year, Windows's overall share of the market will keep dropping.

How bad will it get? According to the latest figures from [Gartner](#), Windows will account for 18.6 percent of the overall personal computing market in 2014, compared to 44.5 percent for Android and 13.9 percent for Mac/iOS. In 2015, Windows hits 16 percent, with Android at 48 percent and Mac/iOS at 15 percent.

In short, Android will account for almost 50 percent of all personal computing devices sold within the next two years, and while Apple's platforms won't quite overtake Windows, they will come pretty close. Both Windows and Mac/iOS become, in effect, secondary platforms, with Android driving three times the unit sales of either. Sales of Android devices will surpass one billion units in calendar year 2014, an amazing milestone.

Related to this surge in Android devices is the popularity of non-PC devices. Smartphones long ago surpassed PC sales and usage, but tablets and so-called hybrid-PCs—tablets with a clamshell or transforming form factor—are perhaps more obviously directly competitive with traditional, non-touch PCs. In 2013, even a diminished PC market that sold about 300 million PCs outperformed tablets, at 180 million units, and hybrid PCs, at about 17 million units. But that's set to change.

In 2014, Gartner expects hardware makers to sell 263 million tablets, a bit lower than the expected 278 million traditional PCs that will be sold. But when you add in 40 million hybrid PCs, that combined market surges ahead of the traditional PC. And in 2015, even standard tablets—which some call media tablets—will outsell the PC.

I've often pegged the hybrid PC as the future of the PC, and here Microsoft and Windows do make a relatively strong showing. But hybrid PCs remain a tiny portion of the overall market—64 million



units in 2015, or about 2.4 percent of all personal computing devices—and even if you lump this device type in with traditional PCs, there’s not much to celebrate. Post-PC it is.

## Microsoft’s Role in the Post-PC World

So what’s Microsoft’s role in this post-PC world? Naturally, the software giant is keen to push Windows as a viable platform going forward. While Windows 8 didn’t get off to a great start, Windows 8.1 was better received, and Microsoft has plans to continue improving Windows over the next few years and, perhaps as important, consolidate the platforms it currently offers for PCs, tablets, and smartphones.

Currently, Microsoft creates three mainstream Windows versions for personal computing devices, and the strategy, and how these products interrelate, is a bit of a mess. The firm offers Windows 8.1 for x86-type traditional PC products, Windows RT 8.1 for ARM-based PCs and tablets, and Windows Phone 8 (also ARM based) for smartphones. But thanks to the recent Microsoft reorganization, that’s about to change.

Terry Myerson, who previously oversaw Windows Phone development and is now in charge of Microsoft’s core and client OS efforts, has indicated that he would like to consolidate Windows RT and Windows Phone into a single platform. That won’t happen overnight, and the change is complicated by the fact that Windows RT and Windows Phone are currently at different stages of development. But it’s certainly possible for Microsoft to move these releases closer to each other over a series of minor updates this year, given their common core, and perhaps to a single platform in time for Windows 9.

Windows 9, you ask? Currently scheduled for an April 2015 release, Windows 9 is code-named Threshold and will be publicly revealed at the April 2014 BUILD Conference (which will otherwise be very much focused on Windows Phone and Xbox).

You can find out more in [“‘Threshold’ to be Called Windows 9, Ship in April 2015.”](#) But before that happens, Microsoft needs to bring

Windows RT and Windows Phone closer together, and it will do so via a series of updates to each this year.

The first, Windows Phone 8.1, has been in development for the past year or more and will bring the Windows Phone platform up to speed with the changes Microsoft made to its other Windows clients last year. A Windows 8.1 Update 1 release, expected in April alongside Windows Phone 8.1, will further improve Microsoft's PC- and tablet-based offerings. And it's possible that future updates (such as Update 2) could target both platforms.

But Microsoft's move to become a devices-and-services powerhouse requires it to focus on more than just Windows. In-house, Microsoft businesses such as Office are already bigger than Windows. And as we discussed earlier, there are bigger platforms outside of Microsoft as well. The firm intends to leave no stone unturned.

This strategy can already be seen in much smaller and consumer-focused offerings, such as Xbox Music—available currently on everything from the web, Xbox One, iPhone and iPod Touch, and Android smartphones and mini-tablets, in addition to Windows—Xbox Live games, Bing, Skype, SkyDrive, and more, across as many popular platforms as possible. If Microsoft makes a service, it will be available on other popular devices via a native mobile app or the web, and not just on Windows.

Of course, Microsoft's biggest strength is in productivity, not the consumer market. And it already makes business-oriented offerings such as OneNote, Lync, SkyDrive Pro, SharePoint, and Office Mobile available on rival platforms, and its [Office 365](#) email, calendar, and contacts services are broadly available on any modern device through Exchange ActiveSync (EAS). These smaller offerings will lead the way to a much bigger opportunity.

In 2014, Microsoft will deliver a “Modern” (what we used to call “Metro”) mobile version of its Office productivity suite for multi-touch Windows tablets and other devices. But the bigger news, perhaps, is that it will be followed up by similar Office releases for iPad

and Android tablets, and possibly quite quickly. I'd be surprised if this year passed without all three releases, in fact.

## What About Us?

If you're focused on Microsoft technologies, you don't need to be told the world is changing, and that's true regardless of what's happening on the client. Despite concerns about privacy and government snooping—some fanciful, some legitimate—2013 was the year that cloud computing took off in a meaningful way, and that trend will only continue going forward.

For all the changes, Microsoft's server- and services-side advantages for businesses of all sizes remain in place. Management of core services—Exchange/Office 365-based email, calendar and contacts, SharePoint-based document management, and so on—won't fundamentally change, though more and more of this infrastructure will be cloud-based rather than in-house.

On that note, Microsoft has done a good job of making its traditional on-premises server offerings available as cloud services, and of course the firm's hybrid deployment offerings—which let businesses move to the cloud at their own speed—are a strength. But even those who stick with Microsoft across the stack will find the ground shifting under their feet.

Business acceptance of consumer-oriented, public cloud services—Dropbox, Google Drive, SkyDrive, and others—will continue to be an issue. In the past, businesses worried about employees walking out the door with vital information on a physical device such as a USB stick or laptop, and then losing that device or having it get stolen. But today, the concerns are broader, and it's more likely for internal information to make its way out of a company, purposefully or not, via a public cloud service.

The changes we see in the personal computing devices space will also continue to wreak havoc with traditional PC management models going forward. If you've not invested in mobile device management



capabilities, this should be the year, as the coming generation of knowledge workers will expect choices when it comes to the devices and services with which they interact. This more agnostic future will include a preponderance of Android devices, iPhones and iPads, Windows PCs and hybrid devices, and even the occasional Chromebook, though Gartner's numbers suggest Google's other computing platform will remain an also-ran for the foreseeable future.

The time to prepare for that future is now. Because it's not the future. It is *now*.

## Embracing Change

Change is something you can either embrace or fear, but it's fair to say that the advent of this devices-and-services era—which is really just another way of saying “post-PC world”—is the biggest change to rock the computing industry since the appearance of the original PC. In some ways, this change is made all the more difficult because previous changes—the rise of Windows, workgroup computing, and the Internet era—were such natural evolutions.

The post-PC world isn't so much a reaction to the past as it is a repudiation of the past, an attempt to move personal computing beyond its complex and old-fashioned beginnings and into a simpler place. We're shocked by this because previous shifts were so less dramatic.

The change from Exchange Server 5.0 to Exchange 2000 Server, for example, was difficult, but not as difficult as the move to a cloud-based Exchange. But what happens when you're asked to move *beyond* Exchange? It's the uncertainty that creates such fear.

And that's why 2014 is going to be a pivotal year for Microsoft, its customers, and those who have bet their careers on this company's technologies. Change is coming. ■

# From One-Liner to ForEach One-Liner

## Get ready for better Active Directory PowerShell tools



**Mark Minasi**

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.

Email



Twitter



Website



In “Where-Object and the Pipeline” and “Introducing the Pipeline and ForEach,” I discussed the two essential building blocks for one-liners: the pipeline and ForEach. This month, I want to show you how to put them together to start building some of those one-liners.

So far, this column’s one-liners have looked like this one, which finds everyone who hasn’t logged on in the past 90 days and disables their accounts:

```
search-adaccount -UsersOnly -AccountInactive -TimeSpan "90" |
  disable-adaccount
```

The first cmdlet, *search-adaccount*, finds the delinquent users and passes them to the second cmdlet, *disable-adaccount*, which disables the accounts; the pipeline performs the passing. The fact that PowerShell lets you do something so useful with such a simple line of text is wonderful because it masks a fair amount of complexity: no looping commands, no variables, and a minimum of punctuation. Unfortunately, if you want to go beyond that fairly basic one-liner and squeeze more automation power out of PowerShell, you need to take a step back. You need to consider how you might have accomplished that *disable-inactive-accounts* task in a way that’s uglier, yes, but that ultimately proves far more flexible.

In particular, consider the second cmdlet, *disable-adaccount*. Typically, you’d have to invoke it with the name of the account that you

want to disable. Open a PowerShell prompt, type just *disable-adaccount*, and press Enter, and PowerShell will complain that you haven't entered an "identity." Add a logon name like this, however, and it would run without error:

```
disable-adaccount JFrost
```

In the one-liner at the start of this column, however, no name was offered to *disable-adaccount*, and it didn't complain; that's because the cmdlet's author built some intelligence into it. If it's invoked without an account name, it looks in the pipeline for anything that looks like an account name, a SID, an object GUID, or a distinguished name (DN). As the *search-adaccount* cmdlet emits entire user objects—all of which have a name, DN, SID, and GUID—a *disable-adaccount* on the receiving end of a pipeline has more than enough information to identify the account to disable. Again, though, it'll accept less, as you can see by just stuffing a samaccountname or a SID into the pipeline and invoking *disable-adaccount* on the other side, like this:

```
"S-1-5-21-3471743624-104184042-3101405554-3104" |  
  disable-adaccount
```

or this:

```
"jdorn" | disable-adaccount
```

In both cases, I've merely shoved some text (a *string*, in programming terminology) into the pipeline. PowerShell then wakes up *disable-adaccount* and drops the pipeline contents into *disable-adaccount*'s lap, and then *disable-adaccount* tries to match the pipeline contents to any known logon names, SIDs, GUIDs, or DNs. If there's a match, great; if not, *disable-adaccount* just displays an error, as it would if I'd typed



```
"roses are red" | disable-adaccount
```

That's not going to match anything that looks like an AD account identifier. Fortunately, most Active Directory (AD)-related cmdlets work the same way; therefore, doing the same thing with *remove-aduser*, *unlock-adaccount*, *enable-adaccount*, and *get-aduser* would succeed as well. (Think twice about experimenting with *remove-aduser*, though!)

Now, let's reconstruct that one-liner, but this time we'll employ *ForEach* and the variable that contains the pipeline's current contents, `$_`. As you've seen, to use *ForEach*, you first fill the pipeline, then insert the pipeline (`()`) and a *foreach*, and then put some commands between the caret brackets—or, in PowerShell terms, a scriptblock. Thus, our one-liner in *ForEach* terms would look something like

```
search-adaccount -UsersOnly -AccountInactive -TimeSpan "90" |  
    foreach {disable-adaccount}
```

That won't work, though, because *disable-adaccount* doesn't automatically get pipeline input when invoked in a *ForEach* scriptblock, so you have to explicitly tell *disable-adaccount* which account to disable. Otherwise, it'll stop dead as it would if you just typed *disable-adaccount* all by itself on a command line. This small change will satisfy *disable-adaccount*:

```
search-adaccount -UsersOnly -AccountInactive -TimeSpan "90" |  
    foreach {disable-adaccount $_ }
```

The `$_` variable satisfies *disable-adaccount*'s need for an account to disable, because at that moment the pipeline contains an AD user object. Where's the percentage in the more complex syntax? More opportunity. For example, that one-liner produces no output at all unless there's an error, so what if you want a log of the accounts you

just deleted? A cmdlet named *add-content* will write a line of text to an existing file if it's followed by a filename and some text, as in

```
add-content c:\logs\disable.log "Disabled another one!"
```

A semicolon lets you put more than one command in a scriptblock, so this modification would log a line for each disabled user's logon name, which is stored in `$_samaccountname`:

```
search-adaccount -UsersOnly -AccountInactive -TimeSpan "90" |  
    foreach {disable-adaccount $_; add-content  
        c:\logs\disable.log $_.samaccountname}
```

Now we're getting closer to building some pretty neat stuff. See you next month for more! ■



**FREE Newsletters!**

**Not your average Newsletters!**  
subscribe today at [windowsitpro.com/manage-newsletters](http://windowsitpro.com/manage-newsletters)

<p><b>WinInfo Daily UPDATE</b> Paul Thurrott covers the entire Windows universe with reviews, commentary, analysis, and tips. Delivered daily.</p>	<p><b>Security UPDATE</b> Learn about Windows security risks, attacks, and how to fix or avoid them. Includes security alerts! Delivered bi-weekly.</p>
<p><b>Windows IT Pro UPDATE</b> Windows industry news, products, FAQs, tips, and resources for IT professionals. Delivered weekly.</p>	<p><b>Dev Pro UPDATE</b> Topics for Microsoft platform developers: ASP.NET, .NET Framework, Silverlight, mobile, and SQL Server development. Delivered weekly.</p>
<p><b>Cloud &amp; Virtualization UPDATE</b> Get the latest news, blogs and analysis to help you determine your organization's cloud and virtualization strategy. Delivered weekly.</p>	<p><b>SQL Server Pro UPDATE</b> The latest news, products, and developments for SQL Server DBAs and developers. Delivered weekly.</p>
<p><b>Exchange and Outlook UPDATE</b> News, strategies, products, and developments in Exchange Server and Outlook messaging. Delivered weekly.</p>	<p><b>SharePoint Pro UPDATE</b> SharePoint for IT professionals and developers – weekly tips, news, and how-to's. Delivered weekly.</p>

**Windows IT Pro**

# Make Windows 8.1 Work Like Windows 7

## Ten tips to recapture everyone's favorite Windows OS



**Michael  
Otey**

is senior technical director  
for *Windows IT Pro* and  
*SQL Server Pro*.

Email



**B**y now the chips are down, and there's no longer any doubt that Windows 8 and its successor, Windows 8.1, have failed to win over most users in the enterprise and consumer markets. The Frankenstein-like mashing of the new Metro touch interface and the traditional Windows desktop is particularly unsettling for typical desktop users who use a keyboard and mouse. Windows 8.1 was Microsoft's rushed attempt to fix the problems in the original Windows 8 release, and although the update helped, it just didn't go far enough. Fortunately, there are several ways you can tweak your Windows 8.1 system to provide a better and more Windows 7-like experience.

### ① **Configure Boot to Desktop**

One of the things desktop users have little use for is the new Metro Start screen; it just gets in the way of accessing the Windows desktop. Fortunately, Windows 8.1 allows you to boot directly to the desktop. To configure this option, right-click the taskbar and select Properties to display the *Taskbar and Navigation properties* dialog box. From the *Taskbar and Navigation properties* dialog box, select *Go to the desktop instead of Start when I sign in*.

### ② **Use the New Start Button**

Removing the Start button and menu from Windows 8 was undoubtedly the worst design decision that Microsoft could possibly have made. Windows 8.1 doesn't fix this problem, either; it just brings

back the Start button, but with no Start menu. However, the new Start button isn't completely without value. If you right-click the Start button, you'll get a handy context menu that allows you to work with Programs and Features, Power Options, Event Viewer, Device Manager, Network Connections, Disk Management, PowerShell, File Explorer, Control Panel, Shutdown, and more. It's no Start menu, but it's better than Windows 8.

### ③ Use the Keyboard Shortcuts

One of the best ways to navigate the new interface in Windows 8.1 and Windows 8 is by using shortcut keys. Fortunately, most of the previous Windows 7 keyboard shortcuts still work. Some of the handy Windows 8.1 and Windows 8 keyboard shortcuts include Alt + Tab to switch between applications, Alt + F4 to close the current application, the Windows key (Win) to switch between the desktop and Start screen, Win + D to display the desktop, Win + L to lock the desktop, Win + R open the Run dialog box, Ctrl + A to select all, Ctrl + C to copy, Ctrl + V to paste, Ctrl + X to cut, and Ctrl + Z to undo.

### ④ Replace the Start Screen with the Apps View

If you're not using Windows 8.1 apps (and honestly, there are very few that have any real use), then you'd probably be better off replacing the Start screen with the Apps view. The Apps view shows you a list of all installed applications, and it doesn't show the Start screen tiles. To enable the Apps view, open the *Taskbar and Navigation properties* dialog box and select *Show the Apps view automatically when I go to Start*.

### ⑤ Show Desktop Background on the Start Screen

If you haven't made a zillion shortcuts on your desktop for all your applications, you'll probably wind up using the Start screen from time to time. If you do, it's nice if it doesn't look completely foreign to, and separate from, the desktop. You can put the Windows 8.1

desktop background on the Start screen view by opening the *Taskbar and Navigation properties* dialog box and then selecting *Show my desktop background on Start*.

---

**Windows 8.1 was  
Microsoft's rushed  
attempt to fix  
Windows 8.**

---

## ⑥ Use the Desktop and Taskbar

Making good use of the desktop and the taskbar are two keys to being productive with Windows 8.1 in a desktop (i.e., keyboard and mouse) environment. Using the taskbar is pretty straightforward. From either the Start screen or the Apps view, you can select an item and choose *Pin to the Taskbar* from the popup menu. Creating desktop shortcuts is a bit more difficult. On the Start screen, click the arrow that appears when you move the cursor (Help desks must love all these invisible options) to display the Apps view. From the Apps view, select the items you want to create shortcuts for and then select *File Locations* from the popup menu. Right-click the desired items, and choose *Send to* and then *Desktop* from the context menu.

## ⑦ Restore Libraries to File Explorer

Another handy Windows 7 feature that Microsoft unceremoniously removed in Windows 8 was the Libraries view option in File Explorer. Libraries are a convenient way to group and access common files. To add the Libraries view, open File Explorer from the desktop and then click the View tab in the ribbon. Next, click the *Navigation pane* button and select *Show Libraries*.

## ⑧ Hide the File Explorer Ribbon

Personally, I like the new File Explorer ribbon. It makes tasks such as displaying file extensions and displaying hidden items very easy by using the new View tab on the ribbon. However, the ribbon is different and does take up window real estate. Unfortunately, you can't natively remove it, but you can hide it by clicking the up arrow in the ribbon's upper right-hand corner.



## 9 Restore the Ability to Play DVDs

Removing the ability to play DVDs was another inexplicable and universally disliked change that Microsoft made to Windows 8, and the Windows 8.1 update does nothing to fix the problem. If you're not using Windows 8.1 Pro, you can download the [Windows 8.1 Pro Pack](#) for \$99.99; if you are using Windows 8.1 Pro, you can buy the [Windows 8.1 Media Center Pack](#) for \$9.99. If you'd rather pay nothing, then you can download the free [VLC media player](#).

## 10 Install a Start Menu Replacement

I can't emphasize this enough, but one of the little things that can really help your Windows 8.1 or Windows 8 experience is to install a third-party Start menu. Why Microsoft didn't just put this back into Windows 8.1 is beyond me. Regardless, [Classic Shell](#) can give you back your Windows 7-like Start menu—and it's free. If you're willing to pay \$4.99, Stardock's [Start8](#) is another great option with a lot of customizable features. Both of these third-party Start menus make Windows 8.1 and Windows 8 a much better desktop experience. ■

# Windows Azure Active Directory Updates in 2013

This baby's had a growth spurt



**Sean Deuby**

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core Directory Services team. He's been a Directory Services MVP since 2004.

Email



Twitter



**T**he Azure Active Directory product group was on a tear in 2013. This multi-tenant directory service for the cloud grew in lock-step with the popularity of Microsoft's online services until the growth made Azure AD the world's largest multi-tenant directory. [As of October 2013, Azure AD adoption](#) included 1.4 million tenants with Microsoft Online Services subscriptions, 240 million user accounts, 10 billion authentication requests in a week, and 430 billion total authentication requests. And the product has been in general availability for only 8 months. Since then, the team has added a substantial list of enhancements to Azure AD to support Microsoft's Cloud OS vision and to bring the cloud directory service into contention with other IDaaS offerings on the market. What was added to Azure AD in 2013, and what might be coming this year? Here's a chronology of Azure AD's enhancements in 2013.

## April

***Azure AD general availability.*** Although it was already supporting hundreds of thousands of tenants and millions of users, [the Azure AD service reached the general availability milestone](#) (a murky definition for web releases).

## June

***Application access enhancements.*** [Azure AD application access enhancements](#) provided single sign-on (SSO) capabilities for SaaS applications. This preview was Azure AD's first clear step into the IDaaS market as a competitor to other IDaaS vendors, such as Centrify, Covisint, OneLogin,

Okta, PingOne, and Symplified. Like these solutions, Azure AD gained an application access panel that presents a graphical list of SaaS applications that users can transparently sign on to. The panel began with fewer than 100 apps but continues to grow on a daily basis.

***Support for multi-factor authentication (MFA) public preview.***

With [support for MFA](#), companies can enable MFA for identities in Azure AD to help secure access to Office 365, Windows Azure, Windows Intune, Dynamics CRM Online, and other apps that are integrated with Azure AD. In addition, you can use the on-premises Active Directory Federation Services (AD FS) role and the Web Application Proxy feature in [Windows Server 2012 R2](#) to create a hybrid MFA solution for both Azure AD users and on-premises Windows Server AD users.

## August

***Public preview of improvements to application access.*** A [public preview of the application access enhancements](#) included peeks at app gallery improvements, improvements in management of password-based SSO apps, and bulk enabling or disabling of MFA for Azure AD identities.

## September

***Create and manage multiple Azure ADs within your Azure subscription.*** Now, [you can manage multiple Azure ADs within your Azure subscription](#). As if AD nomenclature wasn't confusing enough, this enhancement lets you create more than one Azure AD tenant within an Azure subscription. Why would you do this? It's designed for lab or development scenarios, or staging to production purposes. The key word here is *tenant*. In my understanding of the service, there's only one Azure AD overall—and it's a big one. All we can touch are the tenants, which are the individual directory instances.

***MFA general availability.*** Now a billable option, [MFA is charged per user, or per authentication](#). In a boon for security, MFA for Azure AD administrative accounts remains free. Nice touch!

## November

**Azure AD Premium public preview.** Azure AD Premium is Microsoft's first cut at charging for its IDaaS services. The Premium public preview is itself free until this feature set moves to general availability status some time in Q1. Then it will become a billable service (price TBD). The preview includes the following features:

- Self-service password reset (SSPR) for users—Just as its name describes, SSPR allows users to reset their Azure AD account passwords without going through a Help desk. Note that at this time, SSPR is available only to Azure AD accounts.
- Group-based provisioning and access management to SaaS apps—You can use Windows Server AD security groups, copied into Azure AD from Windows Server AD via DirSync (or via Group Management for Admins), to assign access to SaaS apps in bulk. For example, you can create a Salesforce Users security group in Windows Server AD and, once the group synchronizes into Azure AD, assign this group access to your Salesforce app connection.
- Customizable access panel—You can now brand the application access panel with your company logo and colors.
- Machine learning-based security monitoring and reports—This Premium feature promises advanced reporting, including anomalies and inconsistent access patterns, logons by users who logged on from unknown sources, logons that occurred after multiple failures, and logons from multiple geographies in short timespans. It will be interesting to see how these reporting capabilities stack up against those of other reporting tools, and how they evolve.

I'm not sure it's being explicitly branded as such, but the free Azure AD Basic edition offers a subset of the features that are included in the Premium edition. Microsoft has emphasized that the Premium offering's features will continue to grow as the product evolves.

**Application access enhancements general availability.** These changes were moved to general availability status.

***Enhancements to creating and managing multiple Azure AD directories within a subscription.*** The [Azure AD directories enhancement](#) includes improvements such as the ability to easily rename a directory and add users to a new directory from an existing directory.

***Public-facing App Gallery site.*** The [App Gallery update](#) allows potential customers of the Azure AD as an IDaaS service to “window shop” the SaaS apps that the service supports without signing up.

***Enhancements in the GraphAPI.*** The [GraphAPI](#) is the RESTful interface to Azure AD that developers use to extract data from it and to explore the relationships between the data. It’s analogous to using LDAP to query Windows Server AD.

## December

***Group Management for Admins public preview.*** The [Group Management for Admins feature](#) lets you perform create/read/update/delete (CRUD) operations on security groups directly in Azure AD. However, you can assign groups to SaaS applications only if you’re subscribed to Azure AD Premium.

***Custom branding support public preview.*** The [custom branding support feature](#) allows you to create a branded sign-in page. Like the customizable Access Panel, this will be part of the Premium offering when it achieves general availability.

***Open-sourcing Azure AD developer libraries.*** Microsoft is making the [Active Directory Authentication Libraries \(ADAL\)](#) available on [github.com](https://github.com) as open-source libraries.

## And Beyond . . .

What’s does the future hold? I think you can expect greater integration between on-premises Windows Server AD and Azure AD, such as providing SSPR for Windows Server AD via Azure AD. This and other on-premises/cloud-integration capabilities will require a more complex two-way identity flow within the hybrid Windows Server AD/Azure AD identity infrastructure, not just one way from Windows



Server AD to Azure AD. Using AD terminology, this means that either Windows Server AD or Azure AD will be able to perform originating writes to the overall hybrid directory of Windows Server AD plus Azure AD. These updates will then replicate throughout the hybrid identity infrastructure.

If 2013 is any indication, Azure AD will continue to advance strongly in 2014. The cloud directory had a lot of catching up to do compared with its mature Windows Server AD partner's capabilities, and it has definitely closed the gap. But to be clear, I believe Microsoft has no intention of giving Azure AD feature parity with Windows Server AD. Rather, it's building out Azure AD to be a full partner with Windows Server AD to create a hybrid identity infrastructure that supports a hybrid enterprise. The two-way Windows Server AD/Azure AD replication I described would be just one of these supporting changes. ■

# Six Common Problems with Failover Clusters

## Why the problems occur and how to fix them

**A**s you might know, I'm a part of the Microsoft group that supports failover clusters. As a result, I've had to troubleshoot quite a few problems. I'll go over some of the common problems I've seen, explain why they occur, and show you how to fix them.

### Common Problem 1

When the Cluster Service starts, it detects the networks on a node, then identifies the network cards in each network. A common problem that I've encountered is that people are unaware that Windows Server Failover Clustering (WSFC) allows only one network card on a node in the same network. All other cards in that network will be ignored.

For example, suppose an administrator, Bill, configured a node with two cards in the same network:

Card1

IP Address: 10.10.10.1

Subnet Mask: 255.0.0.0

Card2

IP Address: 10.10.10.2

Subnet Mask: 255.0.0.0

The Cluster Network Driver (Netft.sys) will use only one network card (or team) per network. So, in the case of this configuration, Card1 will be used by Cluster Network 1 (10.10.10.0/16) and Card2 will be



**John Marlin**

is a senior support escalation engineer in Windows Commercial Technical Support, focusing on failover clustering. He is a Microsoft Certified Trainer for clustering, delivering to Microsoft and its partners, and is a regular contributor to the [Ask the Core Team](#) blog. He is also a contributor to the new book *Introducing Windows Server 2012* (Microsoft Press).



**Email**



**Website**

ignored by WSFC and not used for any communication between the nodes. Because only one network is being used, if Card1 goes down or loses network connectivity, the node can't communicate with any other nodes. This is a single point of failure. To avoid this problem, you need to configure your cluster so that there are at least two network paths between nodes. That way, if one of the cards goes down, you still have communication between the nodes using the other card.

## Common Problem 2

The second common problem is best described using scenarios. I'll describe the problem using two different cluster configurations: single site and multisite.

**Single-site cluster.** Suppose that Bill the administrator decided to reconfigure his cluster so that it has two networks between Node1 and Node2. On Node1, he changed the network cards' IP addresses and subnet masks to:

Card1

IP Address: 192.168.0.1 (Cluster Network 1)

Subnet Mask: 255.255.255.0

Card2

IP Address: 10.10.10.1 (Cluster Network 2)

Subnet Mask: 255.0.0.0

Bill also changed the IP addresses on Node2 (192.168.0.2 and 10.10.10.2). In addition, on Node1 in the cluster, he added a file server group, giving it the IP address 192.168.0.15.

Afterward, Bill tested the cluster to see whether the file server group would successfully move to Node2 during a failover. However, the IP address failed to come online, so the file server group stayed in an offline state. In the System event log, Bill sees event 1069, with the description that the IP address resource failed.

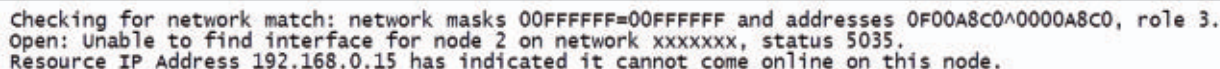
Why did this failure occur? The reason becomes evident if you use the Windows PowerShell `Get-ClusterLog` cmdlet to generate a cluster log. You simply need to run the command:

### Get-ClusterLog

This command will generate a cluster log on each node. To generate a cluster log on only one node, you can add the `-Node` parameter followed by the node's name. You can also add the `-TimeSpan` parameter to create a log that contains information for only the past *x* minutes. For example, the following command generates a cluster log on Node2 that contains information for only the past 15 minutes:

```
Get-ClusterLog -Node Node2 -TimeSpan 15
```

In the results shown in Figure 1, notice “status 5035.” This error message basically tells you that a cluster network isn’t available for the operation. If Bill were to navigate to Networks in Failover Cluster Manager, he’d see that the 192.168.0.0/24 network contains only one network card for Node1. However, there’s a new network 192.0.0.0/8 with Node2’s network card. When Bill changed the network card’s IP address in Node2, he didn’t change the subnet mask. So, error 5035 occurred because Bill misconfigured the card.



```
Checking for network match: network masks 00FFFFFF=00FFFFFF and addresses 0F00A8C0^0000A8C0, role 3.
Open: Unable to find interface for node 2 on network xxxxxxxx, status 5035.
Resource IP Address 192.168.0.15 has indicated it cannot come online on this node.
```

When an IP address resource is created, you have the option to specify the network to use based on the IP address. On its own, WSFC won’t change the network that the IP address resource will use if the network doesn’t exist on the node to which the resource is moving during a failover. In this example, given the IP address specified by Bill and the subnet mask that the IP address will use, the file server group is only going to work on Cluster Network 1 (192.168.0.0/24).

**Figure 1**  
Receiving a Status of 5035 in the Cluster Log File

**Multisite cluster.** In the case of a multisite cluster, each node typically has different IP address networks. When you create the cluster and its roles using the New Resource Wizard, you'll be prompted to enter an IP address for each of the node's networks configured for client access, as Figure 2 shows. When the New Resource Wizard creates the IP addresses and assigns the network name, it automatically gives the network name an "or" dependency. This means that if one of the IP addresses is online, the name will also be online. If you create the groups or resources before adding nodes from a different network, you need to manually create these secondary IP addresses and add the "or" dependency.

**Figure 2**

Creating a Multisite  
Cluster

**New Resource Wizard**

**Client Access Point**

Client Access Point | Confirmation | Configure Client Access Point | Summary

Enter Network Name and IP Address:

Name:

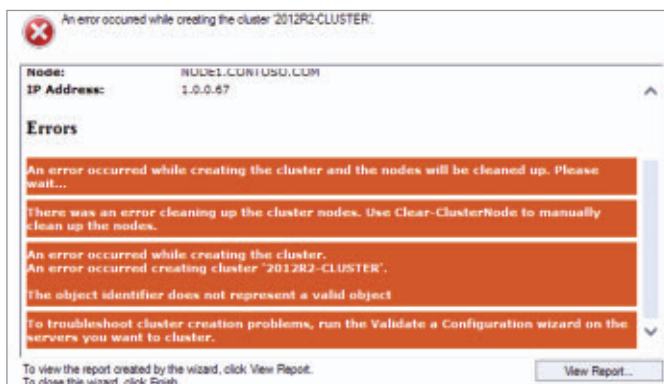
The NetBIOS name is limited to 15 characters. One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.

	Networks	Address
<input checked="" type="checkbox"/>	1.0.0.0/8	Click here to type an address
<input checked="" type="checkbox"/>	192.168.0.0/24	Click here to type an address

### Common Problem 3

When creating a cluster, you don't have to be a domain administrator but you do need to have the proper rights to create objects in [Active Directory](#) (AD). For starters, you need the Read and Create rights on the organizational unit (OU) where the Cluster Name Object (CNO) will be created. The CNO is the computer object associated with a cluster resource called Cluster Name. When creating a cluster, WSFC uses the user account with which you logged on to create the CNO in the same OU where the nodes reside. If you don't have rights to this OU, the cluster creation will fail with the error shown in Figure 3.





**Figure 3**  
Getting an Error When  
Trying to Create a  
Cluster

In “[Troubleshooting Windows Server 2012 Failover Clusters](#),” I mentioned that you can use the Validate a Configuration Wizard in Failover Cluster Manager to help determine the root causes of problems. Using this wizard, you can run numerous tests, including the Validate Active Directory Configuration test. If you run this test and you don’t have rights to the OU, you’ll get errors like those shown in Figure 4. After you fix the rights, you should be able to create the cluster.

### Validate Active Directory Configuration

**Description:** Validate that all the nodes have the same domain, domain role, and organizational unit.

Validating that all nodes have the same domain, domain role, and organizational unit.

Fqdn	Domain	Domain Role	Site Name	Organizational Unit
NODE1.CONTOSO.COM	CONTOSO.COM	Member Server	Default-First-Site-Name	

The distinguished name of node NODE1.CONTOSO.COM could not be determined because of this error: There was an error getting information about the organization unit for node 'NODE1.CONTOSO.COM' from the domain 'CONTOSO.COM'.

The organizational unit of node NODE1.CONTOSO.COM could not be determined because of this error: Did not find an Organization Unit (OU) in the Active Directory

**Figure 4**  
Getting Errors When  
Running the Validate  
Active Directory  
Configuration Test

All other cluster network name resources in the cluster are associated with Virtual Cluster Objects (VCOs), which are created in the same OU as the CNO. Therefore, when creating roles in the cluster, you need to create the CNO with rights (Read and Create) to the OU because the CNO creates all the VCOs in the cluster. If you don’t do so, the new role will be created but the name will be in a failed state. In this case, you’ll see event ID 1194 in the System event log, as shown in Figure 5.

**Figure 5**

Receiving Event ID  
1194 in the System  
Event Log

```
Event ID: 1194
Source: Microsoft-windows-FailoverClustering
Description: Cluster Name resource 'xxxxxx' failed to create its
associated computer object in domain 'yyyyyy' during: Resource Online.
The text for the associated error code is: A constrained violation occurred.
```

There are other settings on the local machine that can cause errors (including access denied errors) when creating VCOs in AD:

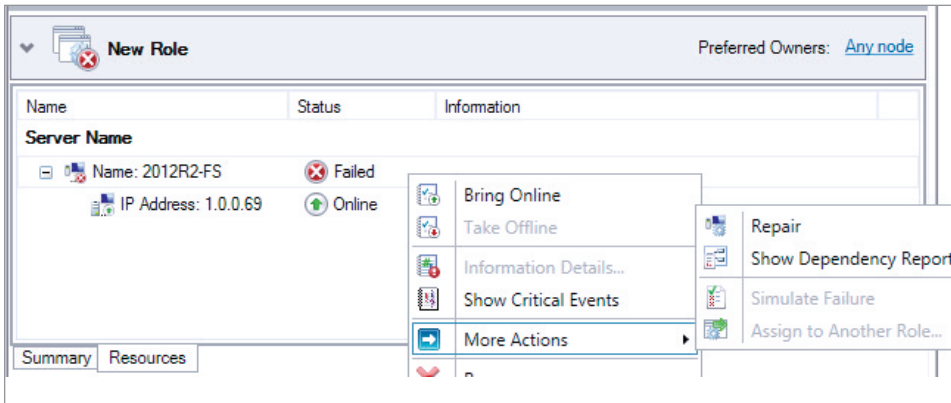
1. The local Users group no longer includes Authenticated Users. It's usually removed by Group Policy Objects (GPOs) or security templates.
2. In the local security policy, the *Access this computer from the network* or *Add workstations to the domain* option no longer includes Authenticated Users. It's usually removed by GPOs or security templates.
3. The following security rights are enabled:
  - *Network access: Do not allow anonymous enumeration of SAM accounts*
  - *Network access: Do not allow anonymous enumeration of SAM accounts and shares*
4. The Cluster Name resource is in a failed state.

## Common Problem 4

The CNO and VCOs are computer accounts—and like user accounts, computer accounts have passwords. AD randomly generates the passwords for computer accounts. By default, the domain policy will reset the password for a computer account every 60 days.

The CNO is used for operations such as joining new nodes to the cluster, creating new objects in the domain, and performing a live migration of virtual machines (VMs) between nodes. The CNO's domain password must be up-to-date for these operations to occur. To be on the safe side, the Cluster Service will attempt to reset the password for its objects at the halfway point (30 days). If the password hasn't been reset at the 60-day mark, the name will fail to come online.

To reset the password, you need to do a repair from within Failover Cluster Manager. As Figure 6 shows, you right-click the failed name resource, select More Actions, and choose Repair.



**Figure 6**

Resetting the CNO Password Manually in Failover Cluster Manager

When issuing a repair, Failover Cluster Manager uses the user account with which you logged on to contact AD to reset the password. Therefore, you must have the Change Password right on the CNO; otherwise, the repair will fail. You also need to make sure that the Reset Password right is enabled on the CNO and VCOs so that WSFC can reset the password when it needs to.

## Common Problem 5

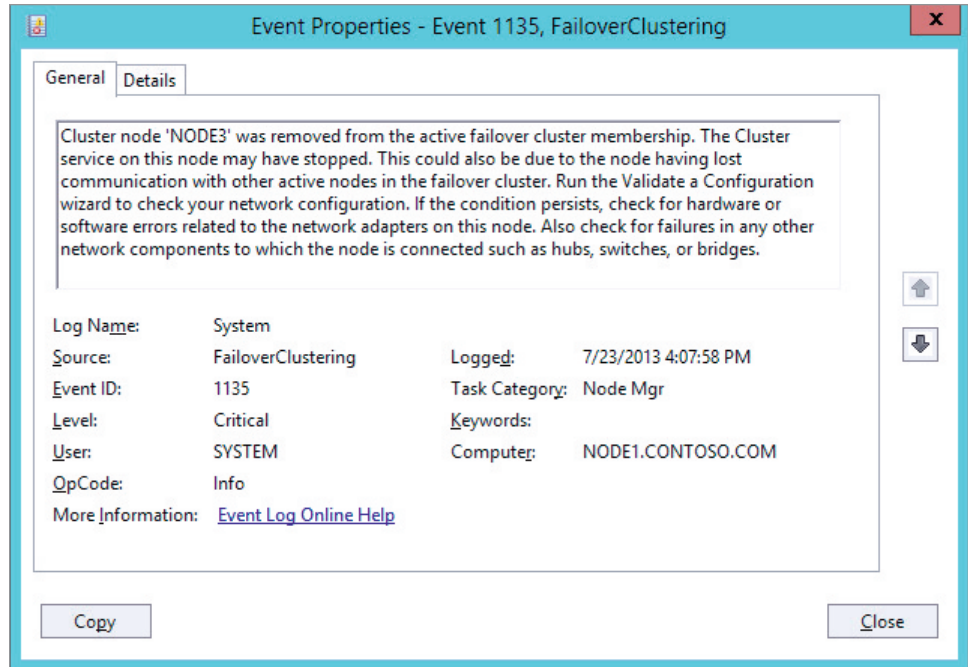
In order for a node to know which other nodes are actively participating in the cluster (i.e., know the current membership), there are a series of heartbeats that go between the nodes over the network. These heartbeat packets are UDP datagrams that travel over port 3343.

Each packet includes a sequence number to track whether the packet is received. Here's how it works: If Node1 sends the sequence number 1111, it expects the return packet to include 1111. This continues between all nodes every second. If Node1 doesn't get the return packet, it will send the next sequence number (1112), and so on.

By default, if the node doesn't receive five heartbeats in five seconds, WSFC determines that the node is down. A participating node

still in the cluster will send a packet to the node determined to be down to terminate the Cluster Service and will log event ID 1135 in the System event log, as Figure 7 shows.

**Figure 7**  
Receiving Event ID  
1135 in the System  
Event Log



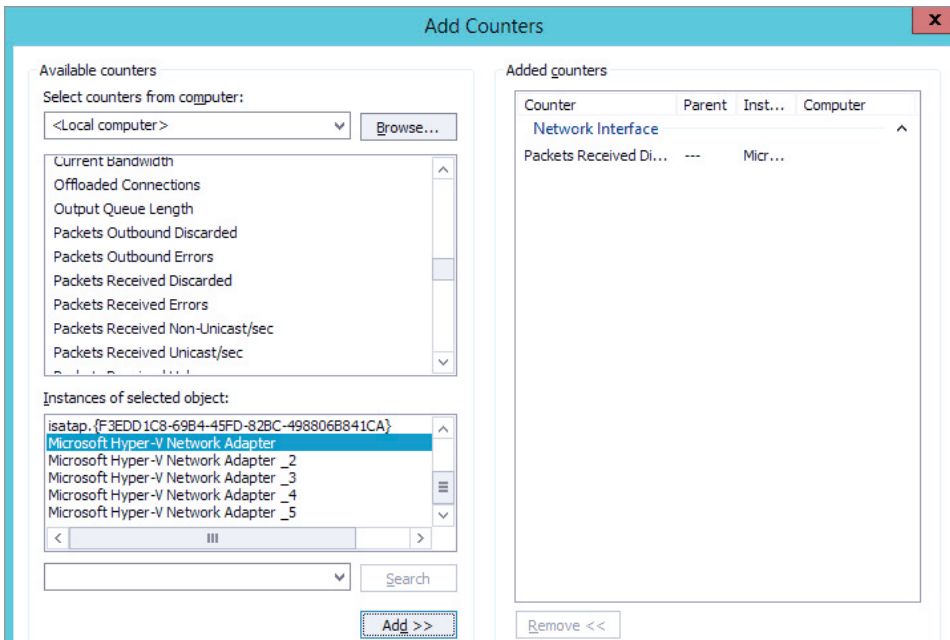
There are multiple reasons why this occurs, many of which involve communication over port 3343 being blocked:

1. Network hardware failures
2. Out-of-date network card drivers or firmware
3. Network latency
4. IPv6 enabled on the servers but the following two rules disabled for inbound and outbound traffic in the Windows Firewall:
  - Core Networking - Neighbor Discovery Advertisement
  - Core Networking - Neighbor Discovery Solicitation
5. Switches, firewalls, or routers not properly configured to allow UDP Datagram traffic
6. Performance problems (e.g., hangs, delays)

## 7. Improperly configured receive buffer settings on the network card driver

One of the first things that I always check is the Packets Received Discarded counter that's part of the Network Interface performance object in Performance Monitor. The Packets Received Discarded counter tracks the number of inbound packets that were chosen to be discarded, even though no errors had been detected to prevent their delivery to a higher layer protocol. The buffer is only so big. If it's not big enough, once the buffer fills up, it must discard packets to make room for additional packets.

To add the Packets Received Discarded counter, open Performance Monitor, right-click its display, and select Add Counters to bring up the Add Counters dialog box. After specifying the appropriate computer, scroll to Network Interface and select the Packets Received Discarded counter. In the *Instances of selected object* drop-down list, choose the appropriate network card and click Add, as Figure 8 shows.



**Figure 8**  
Adding the  
Packets Received  
Discarded Counter in  
Performance Monitor



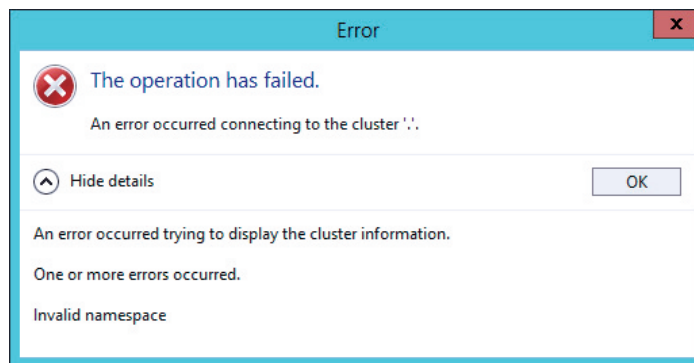
When added, look at the counter's Average, Minimum, and Maximum values. If there are values higher than zero, the receive buffer needs to be adjusted for the network adapter. Check with the vendor of the network card to see what it recommends as a setting. A reboot might be necessary.

In a [Windows Server 2012 R2](#) failover cluster, you can also use the Validate a Configuration Wizard to run the Network/Validate Network Communication test. This test checks to see whether it can communicate between the nodes over port 3343. If it can't, it will post an error and a possible cause.

## Common Problem 6

Sometimes Failover Cluster Manager fails to open, giving you an error message like that shown in Figure 9. When Failover Cluster Manager opens, it opens a Windows Management Instrumentation (WMI) connection to each node in the cluster. In Figure 9, the error message is saying that one of the nodes has an invalid namespace, which means that the Cluster WMI instance (Cluswmi.mof) has been removed from a node. The trick is finding out which node had it removed, because the error message doesn't tell you that information.

**Figure 9**  
Receiving an Error  
Message Noting That  
a Namespace Was  
Invalid



Listing 1 shows a [Windows PowerShell](#) script that you can run to identify the node that's missing the Cluster WMI instance. (You can download this script by clicking the Download button near the top of

### Listing 1: PowerShell Script to Identify Any Nodes Missing the Cluster WMI Instance

```
$NodeNames = Get-ClusterNode
ForEach ($ClusterName in $NodeNames)
{
Write-Host -NoNewline "Testing $ClusterName "
Try
{
$result = (Get-WmiObject -Class "MSCluster_CLUSTER" `
-namespace "root\MSCluster" `
-authentication PacketPrivacy `
-computername $ClusterName -erroraction stop).__SERVER
Write-host " : Successfully queried cluster node "
}
Catch
{
Write-host -NoNewline " : Failed to query cluster node "
Write-host -ForegroundColor Red -BackgroundColor Black `
$_.Exception.Message
}
}
```



Download the code

the page.) After you've identified the node, you can run the following command:

```
Set-Location C:\Windows\System32\Wbem
Mofcomp.exe Cluswmi.mof
```

The most common reason for a node missing Cluswmi.mof actually stems from the old way of fixing WMI. To clear up problems with WMI, administrators would run the command *Mofcomp.exe \*.mof*, which will compile all the Managed Object Format (MOF) files into the WMI repository. The problem is that there are quite a few uninstall files for the various roles and features in Windows, including

Cluster WMI. So when the command is run, it installs Cluswmi.mof, then later uninstalls it. The proper way to rebuild the WMI repository is with the Winmgmt.exe command.

---

**WSFC allows only one network card on a node in the same network. All other cards in that network will be ignored.**

---

## An Ounce of Prevention

As the adage goes, an ounce of prevention is worth a pound of cure. So, I'll conclude by mentioning something you probably already know: You need to keep your machines up-to-date as far as security patches and fixes are concerned. The Microsoft Failover Clustering Team has published articles listing the hotfixes that it would like to see on all clusters. Each Windows version has its own article:

- [“Recommended hotfixes and updates for Windows Server 2012 R2-based failover clusters”](#)
- [“Recommended hotfixes and updates for Windows Server 2012-based failover clusters”](#)
- [“Recommended hotfixes and updates for Windows Server 2008 R2-based server clusters”](#)

These articles are updated as needed, so they're always pretty current. Note that they don't list every fix. Instead, they list the fixes most needed for stability reasons and the most widely requested fixes based on calls coming into Microsoft. ■

# Implementing QoS in Windows Server 2012

Maximizing resource utilization is easier than you think

**T**he famous fictional *Hitchhiker's Guide to the Galaxy* tells readers, first and foremost: "Don't panic." I often think of this message when I try to talk about Quality of Service (QoS) with an organization: Nervous glances quickly fill the room and everyone remembers that they either have a doctor's appointment or a parent/teacher meeting and they need to leave the meeting immediately.

This nervousness isn't entirely without cause. Implementing QoS shouldn't be undertaken lightly. The task involves understanding the types of traffic on the network, the relative importance of that traffic, the network infrastructure itself, and how to actually implement QoS on the network and OSs. QoS is more approachable than ever in *Windows Server 2012*. Organizations have different options for how to implement QoS. But first, why do you even need it?

## The Need for QoS

In today's highly connected data centers, the servers that make up the data center communicate with one another for many reasons:

- Application data
- Replication



## John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Mastering Hyper-V 2012 R2 with System Center and Azure* (Wiley).



- Cluster traffic
- Network storage traffic (e.g., Server Message Block—SMB, iSCSI)
- Management traffic
- Backup data

Add in virtualization and things get even more complicated. On a Hyper-V server, I might need five network connections, minimum! Also consider that redundant paths might be required for some types of traffic, so multiple connections might be teamed together. Separate connectivity might be needed for storage when technologies such as Fibre Channel are used. All this can make a potential mess. So how do organizations architect their fabric?

Think of your network as a highway. You hear a siren behind you. Cars around you begin trying to move out the way, but doing so during rush-hour gridlock is difficult and the emergency vehicle can't get through—a disaster for whomever is relying on the aid it provides. Or perhaps the highway you're on has a special lane for emergency vehicles. Most times, that lane is empty. As you sit in rush-hour traffic, your anger builds as you realize how much quicker you could get home if only other cars could use that lane. Of course, the traffic department could just keep adding lanes to the highway to avoid congestion, but that isn't a practical solution.

Your network is like that highway. Maybe all the network traffic shares a network connection, in which case you risk crucial traffic being unable to traverse the network in a timely fashion during times of high load. Or perhaps you have dedicated network connections for each type of traffic to ensure that it always gets the bandwidth it needs when it needs it. Maybe you keep adding connections to accommodate all the traffic.

Although most organizations have traditionally used the second option, doing so has become far more challenging for several reasons:

- Data centers are shifting to 10Gb networks instead of 1Gb networks. Having more than two 10Gb connections per server isn't

cost-effective, so having dedicated connections for each type of traffic no longer makes sense.

- As virtualization becomes more prevalent, blade-type servers become more common. But these servers typically have limitations on the number of supported adapters, limiting connections. There are some exceptions if the data center uses converged fabrics, which allow virtual adapters to be created for the host, giving almost unlimited flexibility in how the traffic is divided (although behind the scenes, this is really its own kind of QoS).
- Traditionally, networks have been heavily over-built to ensure that bandwidth is available. Because of the increased importance and use of different types of network traffic, this over-building and complexity have become unmanageable for most organizations. Many of the dedicated network connections for a specific type of traffic are either not used or barely used the majority of the time.

Virtualization has also introduced its own challenges. Many OS instances now run on a single piece of hardware and share a set of network connections. Maintaining separate network adapters for every virtual machine (VM) is highly impractical. A single bad VM can consume all available network bandwidth, starving the other VMs. Therefore, you need a mechanism to ensure not only that different types of traffic have sufficient bandwidth when they need it but also that you can fairly divide network resources of the same type for the different VMs, or tenants, that use the virtual infrastructure.

Imagine that you're a hoster and have many tenants using your services. You need to ensure that each tenant gets a fair amount of resources. You also might need different levels of network speeds, such as Gold, Silver, and Bronze.

As you might expect, QoS is the solution to these problems. QoS fully embraces and enables the multi-tenant and converged 10Gb fabric data centers that we're seeing more frequently. Let's look at the types of QoS that are available through Windows Server 2012,



including software-based QoS, hardware-enabled QoS, and QoS that is specific to virtualization.

You might think the problem can be solved just by adding 10Gb instead of 1Gb to your network. After all, surely a pipe that is 10 times larger than what was available previously will be big enough? That's the same analogy as adding more lanes to the highway: For a while, it might cure the problem, but no matter which resources you give to a workload, it will ultimately grow to use them all and want more. Even a 10Gb connection will become saturated over time, and the problem—certain types of traffic not getting enough bandwidth when needed—will return.

## Software QoS

Windows has had the ability to implement software QoS for a long time. Navigate within a Group Policy Object (GPO) to Computer Configuration, Policies, Windows Settings, Policy-based QoS. Policies can be created to throttle bandwidth to a specific speed for different applications, source-and-destination IP address combinations, and specific protocols. But QoS previously specified a *maximum* bandwidth per traffic type as the means to manage the traffic.

With a maximum-bandwidth configuration, the workload that is affected by the policy can never exceed the amount of allocated bandwidth, which guarantees and enforces a predictable network throughput. For example, suppose that I have a 10Gb network connection and I divide the traffic this way:

- 1Gb for management
- 1Gb for live migration
- 1Gb for cluster or cluster shared volume (CSV)
- 2Gb for iSCSI
- 5Gb for VMs

This looks great. Certainly, in normal circumstances, the VM traffic would be the majority, and the other types of traffic are guaranteed

their own amounts for the times they need it. But that's the problem with using maximum-bandwidth settings for bandwidth management: The majority of the time, 50 percent of the available network bandwidth is hardly used, but it's been reserved for the times that it is needed, just like having a separate lane for emergency traffic. Meanwhile, the servers have a heavy virtualization traffic load and could likely benefit from using that 50 percent of the bandwidth when it isn't in use. We're wasting a large amount of our resources and restricting our capabilities unnecessarily. (Furthermore, QoS in Windows Server 2008 R2 didn't work with Hyper-V.)

One area in which the maximum-bandwidth method is useful is when you must pay for bandwidth usage, such as a WAN connection between offices. In such a scenario, limiting bandwidth to a specific value is a good idea.

Windows Server 2012 introduces the *minimum*-bandwidth policy concept, which allows the different types of traffic to have a relative weight assigned to them. Let's see how that would work for the same types of traffic that we used in the maximum-bandwidth example:

- 10 for management
- 20 for live migration
- 20 for cluster or CSV
- 10 for iSCSI
- 40 for VMs

Note that these values don't represent any kind of unit; they're just relative weights.

The way that minimum bandwidth works is that by default, any traffic type can use any available network bandwidth. Even though VM traffic has a weight of 40, that traffic could consume 100 percent of the network bandwidth as long as there is no contention on the network. When there is no contention, workloads can use whatever bandwidth they want, up to the limit of the network fabric itself. Only in times of contention do the minimum-bandwidth relative weights

come into play. In times of contention, the different types of traffic are guaranteed to get bandwidth based on these weights: Live migration traffic would get 20 percent, while VM traffic would get 40 percent. All the weights together add up to 100. The actual minimum bandwidth for one type of traffic in contention scenarios can be found by dividing the relative weight of that traffic type by the sum of all the weights.

It is also possible to use strict minimum-bandwidth configuration, in which the types of traffic are given absolute bandwidth values as the minimum. For example, management traffic is given 1Gb and VM traffic is given 4Gb. However, this approach can be difficult to administer, compared to relative weightings. You must be careful not to over-provision the minimum values that you assign to the workloads. Also, just because a type of traffic is set to a minimum of 1Gb does not guarantee that it will get 1Gb. Most networks have many switches and routes in them. Although a type of traffic might be guaranteed 1Gb of bandwidth within the local server, after that traffic leaves the server it is at the mercy of the other traffic on the network.

There is another problem with using strict minimum bandwidth. Suppose that you use NIC teaming with two 10Gb network adapters; 20Gb of bandwidth is available when the environment is healthy. If you use strict minimum bandwidth and configure up to 20Gb of minimum bandwidth, those bandwidth guarantees cannot be met if an adapter fails. That defeats the whole point of NIC teaming: to provide consistent service even when a failure occurs.

For these reasons, the strict minimum-bandwidth approach is not recommended. Use relative weighting whenever possible.

The biggest difference between using a maximum- or minimum-bandwidth approach is that the minimum-bandwidth option allows the highest utilization of network resources when there is no contention, whereas the maximum-bandwidth option always limits traffic to its configured maximum, potentially wasting available bandwidth. There is generally no price for intra-data center bandwidth, so using

as much as possible is ideal. If you are using QoS from before Windows Server 2012, then migrating to the minimum bandwidth policies allows you to maximize your resource utilization and avoid waste.

## Hardware QoS

In addition to software QoS, modern networking equipment also enables its own QoS capabilities. This type of hardware-level QoS is known as Data Center Bridging (DCB) and is an IEEE standard. This article focuses more on software-level QoS, but it's important to understand that if you have DCB-enabled network hardware, then you can leverage those capabilities through Windows Server 2012. Traffic management is offloaded to the network adapter instead of being processed by the host OS. In addition, DCB can support flow control for certain types of network traffic and can request a slowdown of traffic from the source (typically a switch). The good news is that most 10Gb network equipment actually supports DCB. By default, DCB is not enabled in Windows Server 2012 and must be installed as a built-in feature, using the following command:

```
Install-WindowsFeature Data-Center-Bridging
```

## Virtualization QoS

Software QoS capabilities are also available to VMs. The ability to have QoS with VMs is crucial for many environments, especially hosters or organizations that have different business units sharing the infrastructure. Being able to ensure that different tenants get fair amounts of network resources and to enable different levels of service based on premium rates for a Gold network connection are huge benefits. Virtualization offers these capabilities for CPU, memory, and even storage, so being able to offer them for the network completes the resource-management picture.

Remember that minimum-bandwidth capabilities can use relative weighting or a strict bandwidth value. With VMs, using minimum

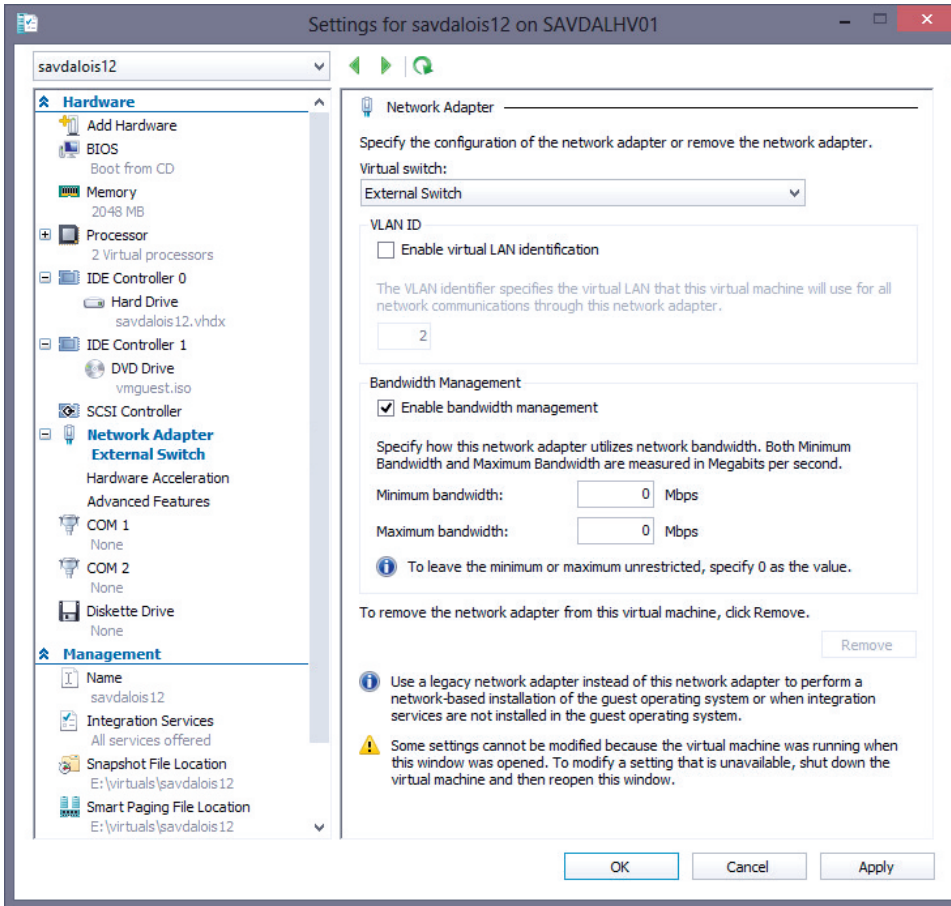
relative weighting is even more important because VMs are mobile. A VM can be moved between hosts, so trying to set strict bandwidth minimums won't work because different hosts have different VMs with their own configurations. If you try to use live migration to move a VM to another host on which the strict minimum-bandwidth configuration cannot be satisfied, then the live migration will fail. Using relative weighting always works because the bandwidth is relative to whichever workloads are on a server.

It is important to understand that minimum-bandwidth QoS policies applied to VMs affect only traffic that is sent from the VM to the physical wire. If the traffic is between VMs on the same host, then minimum-bandwidth QoS policies do not apply. VM-to-VM traffic on the same host never touches the physical network adapter but is routed internally by the Hyper-V virtual switch and therefore does not use up any bandwidth on the network. Maximum-bandwidth QoS policies *do* apply for VM-to-VM traffic, in addition to VM-to-wire traffic. The reason for the difference is that some organizations charge based on network resources used. If a maximum amount of bandwidth is set, then tenants do not expect to be charged more than the maximum that they have configured.

Note that this maximum bandwidth applies only to outbound traffic from the VM (egress); inbound (ingress) is unrestricted. Why not limit inbound traffic? Inbound traffic is already at the host, so dropping it would not bring much benefit. Also, unless the traffic is TCP, there is no way to tell the sender to slow down or stop.

For VMs, only strict bandwidth can currently be configured using the Hyper-V Manager GUI, as shown in Figure 1. This is unfortunate, given that best practice is *not* to use strict minimum bandwidth but rather to use relative weight minimum bandwidth.

The good news is that like the rest of Windows Server 2012, everything that can be accomplished with the GUI can also be accomplished using Windows PowerShell, which actually exposes more functionality, including the relative weight configurations. In addition, through



**Figure 1**  
Hyper-V Manager GUI  
Restrictions

PowerShell you can configure QoS on the Hyper-V virtual switches—something that is impossible using the graphical tools.

## Managing QoS

Many QoS configurations are performed using PowerShell, and the management is consistent between software and hardware QoS or DCB for the majority of actions. The first task is to classify the types of network traffic so that policies can be applied.

When using hardware QoS, there is a limit of eight classifications of traffic; there is no such limitation when using software QoS. For



example, you might create an iSCSI type classification and a Live Migration type classification. The good news is that the PowerShell modules that are used to create the classifications have a number of built-in classifications that include the most common types of traffic (i.e., iSCSI, NFS, SMB, Live Migration, SMB Directory, and Wild Card, which covers everything else). If you need other classifications, you can create your own filters.

After the data is classified, you can create and apply policies to control the allocated bandwidth. Rather than list the cmdlets here, I recommend that you review the Microsoft article “[Network Quality of Service \(QoS\) Cmdlets in Windows PowerShell](#),” which goes into detail about each cmdlet that you need to use.

The following is a simple example that creates a new policy for Live Migration type traffic, using the built-in Live Migration filter:

```
New-NetQosPolicy "Live Migration" -LiveMigration
-MinBandwidthWeightAction 40
```

Another option for using QoS relates to a Hyper-V host, because of the virtual switches that are present. Virtual switches are exposed to the management OS, not just to the VMs. Therefore, you can create multiple virtual adapters for use by the Hyper-V host itself. For example, you can create a live migration virtual network adapter, a cluster virtual adapter, and so on. You can then apply QoS policies directly to the virtual network adapters to control their bandwidth. Doing so might be easier than the typical requirement to use all the separate types of filters to differentiate between traffic. For example, the following two commands create a new virtual adapter on the Hyper-V host OS and then assign it a policy of a minimum-bandwidth relative weight:

```
Add-VMNetworkAdapter -ManagementOS -Name "LiveMigration"
-SwitchName "External Switch"
```


```
Set-VMNetworkAdapter -ManagementOS -Name "LiveMigration"  
-MinimumBandwidthWeight 40
```

## Get on the Road

The primary message surrounding Windows Server 2012 QoS is that it can change the way you think about networking in the data center. Instead of having separate network connections for each workload, use larger pipes and use QoS to ensure that different types of traffic receive the necessary bandwidth. Even if you have DCB-capable network hardware, you might want to use software QoS, especially in Hyper-V environments; software QoS is more scalable in terms of number of policies. Some work is involved in implementing QoS, but the new relative minimums make the process much easier while ensuring that you get maximum utilization of resources. ■

# CAN'T GET AWAY?

Get first-class education from your desk




*Windows IT Pro* offers FREE online events including webcasts, demos and virtual conferences. All events are brought to your computer live while being fully interactive.

Go to <http://windowsitpro.com/web-seminars> to see an up-to-date list of all online events.

## Windows IT Pro

First-class education from the top experts in the industry.  
Visit <http://windowsitpro.com/web-seminars> for a knowledge upgrade today!

**Have a full plate on the live date? Don't sweat it! All online events are recorded and available 24/7.**



# Windows IT Pro Store

---

**eLearning Classes**

---

**eBooks**

---

**On-Demand Training**

---

**In-Person Training**

---

**Posters**

---

**Videos**

---

Plus you can **RENEW** your subscription or  
**UPGRADE** to VIP membership while  
you're there!

**Stop by the store today!**

**Windows**ITPro

# Managing Active Directory with PowerShell

## How to use the Active Directory module

When [Windows PowerShell](#) first shipped, one of the most commonly asked questions was, “Can I manage Active Directory (AD) using PowerShell?” At the time, Microsoft’s answer wasn’t what most administrators wanted to hear. PowerShell had a built-in Active Directory Service Interfaces (ADSI) “type accelerator” that let you access AD objects, but you were pretty much on your own figuring out how to make it work to perform AD administrative tasks. Shortly thereafter, Quest Software offered a free set of cmdlets for performing AD administration tasks, such as creating, modifying, and deleting AD objects and searching for objects in AD. For a long time, this was the state of PowerShell and AD management.

When Microsoft shipped Windows Server 2008 R2, everything changed because it introduced the *Active Directory Module for Windows PowerShell*. The AD module includes a set of cmdlets for managing AD as well as an AD Provider that lets you navigate AD as if it were a drive letter. I’ll describe how to install the AD module and how it works in detail.

### Installing the Active Directory Module

Unlike previous tools that use LDAP to communicate with AD, the AD module uses the Active Directory Web Services (ADWS) protocols to communicate with an AD domain controller (DC). The MSDN blog posting “[Active Directory Web Services Overview](#)” describes these communication protocols in detail; both the PowerShell cmdlets in the AD module and the Active Directory Administrative Center (ADAC) use ADWS to communicate with and get information from AD.



### Darren Mar-Elia

is a Microsoft MVP for Group Policy, a contributing editor for *Windows IT Pro*, and CTO and founder of [SDM Software](#). He maintains a Group Policy resource website and has authored many books on Group Policy and Windows topics.



Email



Website

When you install [Windows Server 2012](#) or Server 2008 R2 DCs in your AD domain, ADWS will be installed and running by default on each of them. If you have a domain composed entirely of Windows Server 2008 or Windows Server 2003 DCs, you need to do a separate ADWS install. Microsoft provides the free [Active Directory Management Gateway Service](#) package for this purpose. If you install this package on at least one Server 2008 or Server 2003 AD DC in your domain, you can use the AD module for PowerShell as well as ADAC.

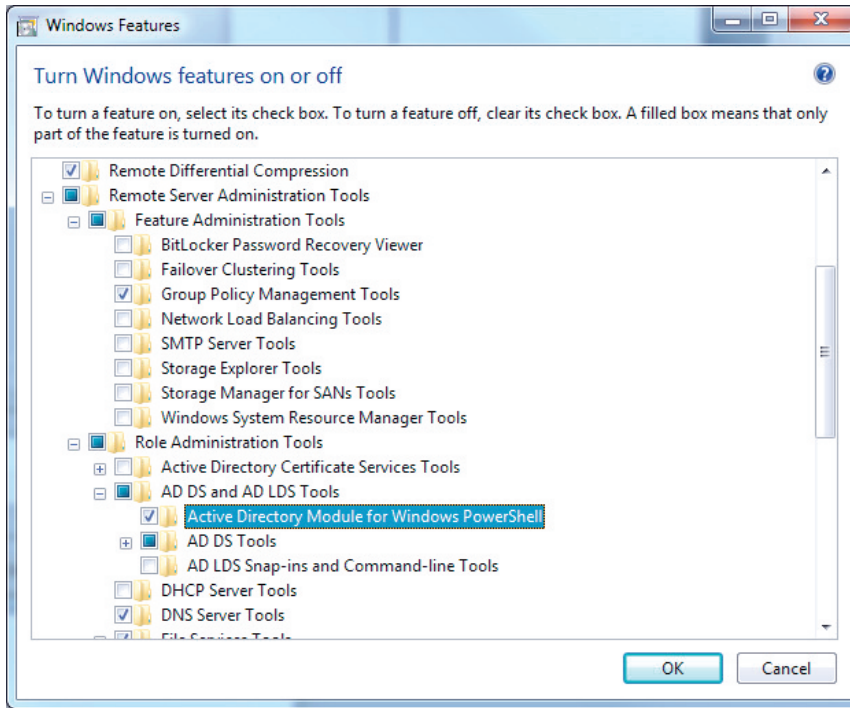
The AD module itself is installed by default on any DC running Server 2012 or Server 2008 R2. If you're running a [Windows 8](#) or Windows 7 box (or any non-DC running Server 2012 or Server 2008 R2), you need to install the Remote Server Administration Tools from the [Microsoft Download Center](#).

No matter whether the Remote Server Administration Tools were already on your system or you installed them separately, the next step is to open the Control Panel Add/Remove Programs applet and select *Turn Windows features on or off* from the menu on the left. In the Windows Feature dialog box that appears, scroll down to the Remote Server Administration Tools section. Look for the *Active Directory Module for Windows PowerShell* check box, which will be in the \Remote Server Administration Tools\Role Administration Tools\AD DS and AD LDS Tools folder, as shown in Figure 1. Select that check box and click OK to install the module.

Afterward, you should see a shortcut labeled *Active Directory Module for Windows PowerShell* under Administrative Tools on the Start menu. Clicking that shortcut will launch PowerShell with the AD module loaded. If you're already working in PowerShell and simply want to load the module so it's available for use, you can type the following command to get access to the AD cmdlets and the AD Provider:

```
Import-Module ActiveDirectory
```

Now let's look at how you can navigate AD using the AD Provider.



**Figure 1**  
Installing the AD  
Module for PowerShell

## Using the Active Directory Provider

PowerShell incorporates the concept of *PowerShell drives*, which I like to simply refer to as PS drives. In simple terms, a PS drive is a way of representing a resource like a navigable file system that's composed of folders and leaf items. Not every resource can be represented this way, but many—including AD and the registry—fit well into that model. The AD module contains the provider for an AD PS drive. What this means is that you can navigate and even modify AD as if it were a file system.

So, how do you navigate AD using the AD Provider? Assuming that you already have PowerShell open and the AD module loaded, the first step is to run the Set-Location cmdlet, which has several aliases, including *sl* and *cd*:

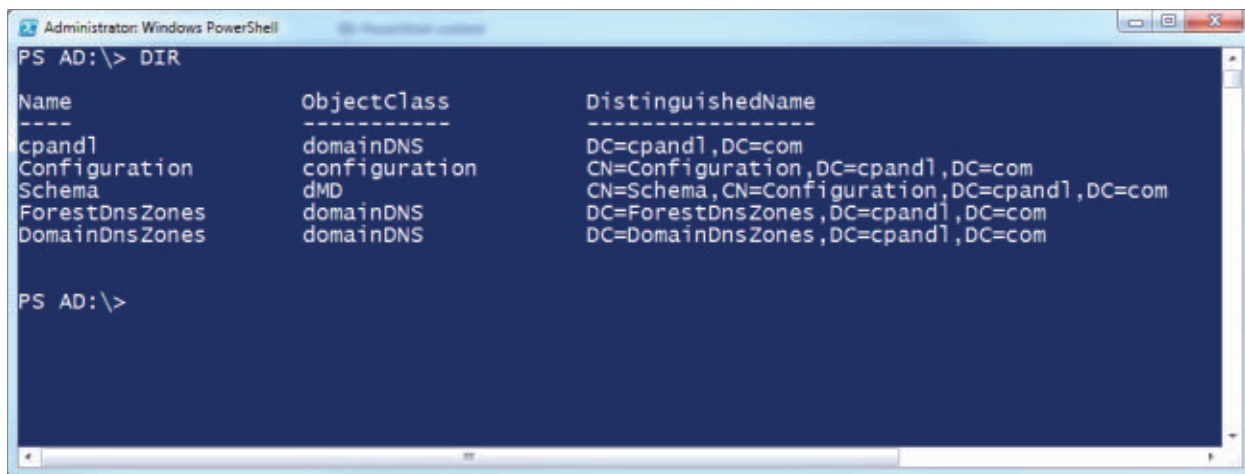
Set-Location AD:



This command changes the current working location to the AD PS drive. As a result, the PowerShell prompt will show *AD:\* instead of *C:\*. Next, to see the items in the AD PS drive, you can use the *Get-ChildItem* cmdlet, which has an alias of *dir*:

*Get-ChildItem*

Figure 2 shows sample results from my machine.



```

PS AD:\> DIR

Name                ObjectClass          DistinguishedName
----                -
cpand1              domainDNS            DC=cpand1,DC=com
Configuration        configuration        CN=Configuration,DC=cpand1,DC=com
Schema              dMD                  CN=Schema,CN=Configuration,DC=cpand1,DC=com
ForestDnsZones      domainDNS            DC=ForestDnsZones,DC=cpand1,DC=com
DomainDnsZones      domainDNS            DC=DomainDnsZones,DC=cpand1,DC=com

PS AD:\>

```

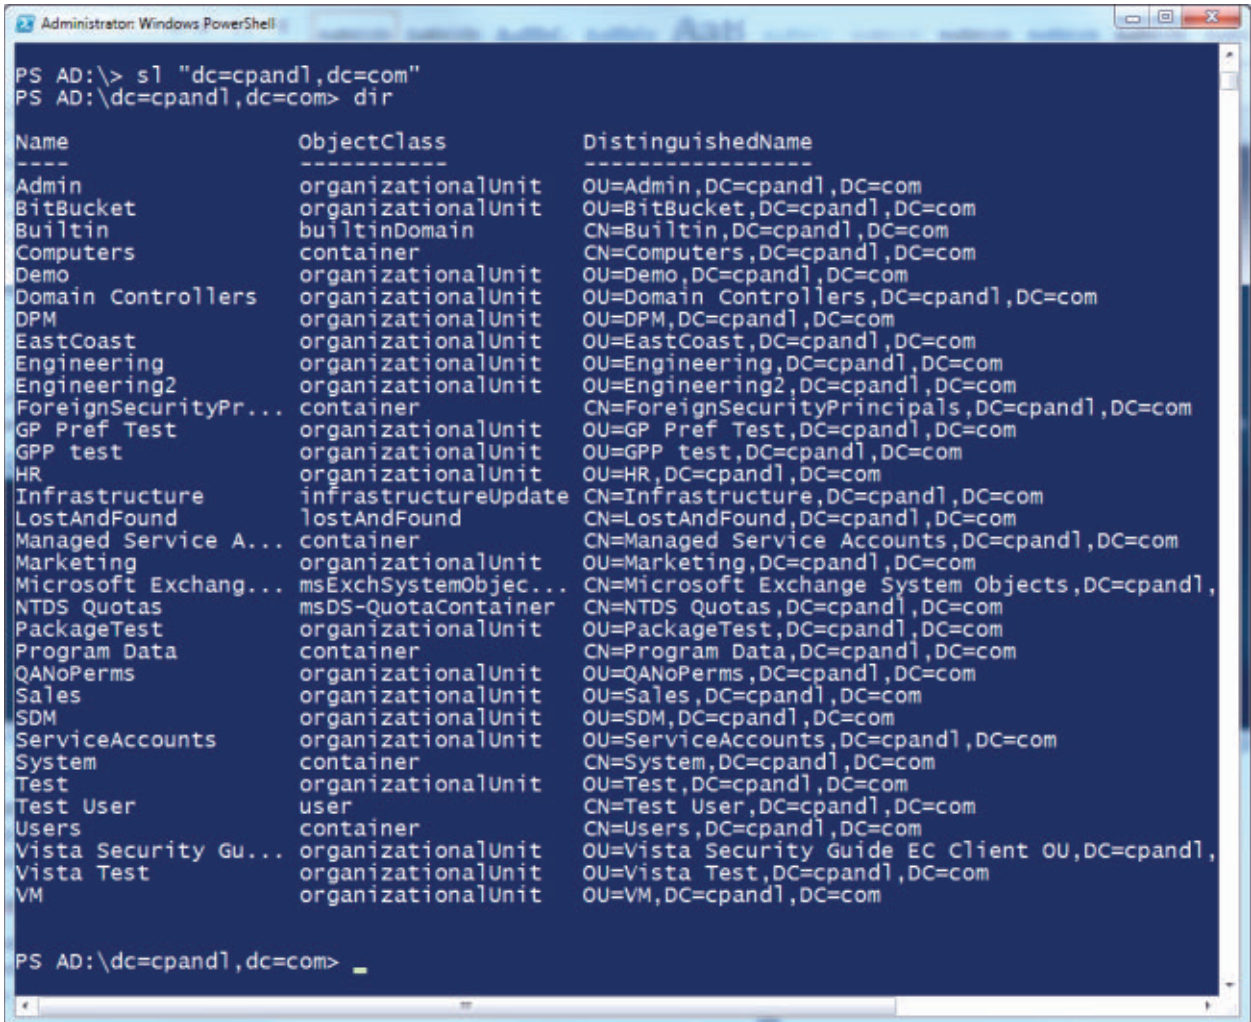
**Figure 2**

Listing the Domain  
Partitions Available in  
the AD PS Drive

As you can see, this command returned a list of all the available domain partitions. The most interesting one for me is the domain partition named *cpand1*, which contains users and computers. To change to that domain, I simply type:

*Set-Location "dc=cpand1,dc=com"*

Note that I'm using the *Set-Location* cmdlet with the distinguished name (DN) of my AD domain. That's required to properly navigate into it. Once I'm in the domain directory (as indicated by *AD:\dc=cpand1,dc=com* in the PowerShell prompt), I can use the *Get-ChildItem* cmdlet to see my top-level AD structure, which Figure 3 shows.



```

PS AD:\> sl "dc=cpan1,dc=com"
PS AD:\dc=cpan1,dc=com> dir

Name                               ObjectClass                DistinguishedName
----                               -
Admin                             organizationalUnit         OU=Admin,DC=cpan1,DC=com
BitBucket                         organizationalUnit         OU=BitBucket,DC=cpan1,DC=com
Builtin                           builtinDomain              CN=Builtin,DC=cpan1,DC=com
Computers                         container                  CN=Computers,DC=cpan1,DC=com
Demo                             organizationalUnit         OU=Demo,DC=cpan1,DC=com
Domain Controllers                organizationalUnit         OU=Domain Controllers,DC=cpan1,DC=com
DPM                              organizationalUnit         OU=DPM,DC=cpan1,DC=com
EastCoast                        organizationalUnit         OU=EastCoast,DC=cpan1,DC=com
Engineering                       organizationalUnit         OU=Engineering,DC=cpan1,DC=com
Engineering2                     organizationalUnit         OU=Engineering2,DC=cpan1,DC=com
ForeignSecurityPr...              container                  CN=ForeignSecurityPrincipals,DC=cpan1,DC=com
GP Pref Test                      organizationalUnit         OU=GP Pref Test,DC=cpan1,DC=com
GPP test                         organizationalUnit         OU=GPP test,DC=cpan1,DC=com
HR                               organizationalUnit         OU=HR,DC=cpan1,DC=com
Infrastructure                    infrastructureUpdate       CN=Infrastructure,DC=cpan1,DC=com
LostAndFound                     lostAndFound              CN=LostAndFound,DC=cpan1,DC=com
Managed Service A...             container                  CN=Managed Service Accounts,DC=cpan1,DC=com
Marketing                         organizationalUnit         OU=Marketing,DC=cpan1,DC=com
Microsoft Exchange...            msExchSystemObjec...      CN=Microsoft Exchange System Objects,DC=cpan1,
NTDS Quotas                      msDS-QuotaContainer       CN=NTDS Quotas,DC=cpan1,DC=com
PackageTest                      organizationalUnit         OU=PackageTest,DC=cpan1,DC=com
Program Data                     container                  CN=Program Data,DC=cpan1,DC=com
QANoPerms                       organizationalUnit         OU=QANoPerms,DC=cpan1,DC=com
Sales                            organizationalUnit         OU=Sales,DC=cpan1,DC=com
SDM                              organizationalUnit         OU=SDM,DC=cpan1,DC=com
ServiceAccounts                  organizationalUnit         OU=ServiceAccounts,DC=cpan1,DC=com
System                           container                  CN=System,DC=cpan1,DC=com
Test                             organizationalUnit         OU=Test,DC=cpan1,DC=com
Test User                        user                      CN=Test User,DC=cpan1,DC=com
Users                            container                  CN=Users,DC=cpan1,DC=com
Vista Security Gu...              organizationalUnit         OU=Vista Security Guide EC Client OU,DC=cpan1,
Vista Test                       organizationalUnit         OU=Vista Test,DC=cpan1,DC=com
VM                               organizationalUnit         OU=VM,DC=cpan1,DC=com

PS AD:\dc=cpan1,dc=com>

```

Suppose I want to look at the users in the SDM organizational unit (OU). To get into that OU, I simply type:

```
Set-Location "OU=SDM"
```

The PowerShell prompt will now show *AD:\ou = SDM,dc = cpan1,dc = com*. At this point, I can use the `Get-ChildItem` cmdlet to see all the user objects in that OU. Let's say I want to change the Description

**Figure 3**

Viewing the Top-Level AD Hierarchy

property on the user object representing my user account *Darren Mar-Elia*. There's a cmdlet for that! The `Set-ItemProperty` cmdlet lets you change a property in an AD object. If I want to change my user account's description to *Chief Techie*, I'd run the command:

```
Set-ItemProperty -Path '.\CN=Darren Mar-Elia' `
  -Name "Description" -Value "Chief Techie"
```

As you can see from this command, I'm using the cmdlet's `-Path` parameter to point to my user account in the current directory. I'm also using the `-Name` parameter to indicate that I want to modify the `Description` property and the `-Value` parameter to indicate that I want the description to be *Chief Techie*.

Note that if you want to find all objects that have a particular property value, you can use the `Get-ItemProperty` cmdlet. If you just want to get a reference to an AD object, the `Get-Item` cmdlet will do the trick.

As you can see, it's pretty straightforward to work with AD this way. Although it might not be a mechanism you'd use for doing mass changes, it's handy to be able to deal with AD as if it were a file system. With that said, I find that most administrators use the AD cmdlets rather than the AD PS drive to manage AD. So, let's see how some of these cmdlets work.

## Using the Active Directory Cmdlets

The AD module that comes with Windows 7 contains 76 cmdlets for managing AD. You can use them for doing pretty much everything, including searching AD objects, creating and deleting AD objects, and manipulating AD configuration information (e.g., forest mode, fine-grained password policy). The cmdlets are generally grouped by their verbs, such as `Add-`, `Remove-`, `Get-`, and `Set-`. Note that not every `Get-` cmdlet includes a corresponding `Set-` cmdlet and vice versa, so you might have to do some digging to find the cmdlet that's right for a task. For example, you can set the AD forest functionality level by using the

Set-ADForestModecmdlet, but if you want to find out the current forest functionality level of a forest, you need to use the Get-ADForest cmdlet and view the ForestMode property on the returned object.

Now let's take a look at some common tasks that you can perform using the AD cmdlets. Specifically, I'll show you how to add user accounts, manage group membership, reset user account passwords, and search for AD objects.

## Adding User Accounts

The New-ADUser cmdlet provides an easy way to add user accounts to AD. Suppose I want to add a new user account named Bill Smith to my SDM OU. In the most basic form, I can create a new user using the command:

```
New-ADUser -Name "Bill Smith" -SamAccountName "bsmith" `
  -GivenName "Bill" -Surname "Smith" `
  -DisplayName "Bill Smith" -Path "OU=SDM,DC=cpandl,DC=com"
```

In this command, I'm filling in some basic information about the user account. Most notably, I'm using the -SamAccountName parameter to provide the SAM account name, which is required to create a user object. I'm also using the -Path parameter to tell the cmdlet where to put the object—in this case, in my SDM OU in the cpandl.com domain. In addition, I'm providing the user's first name (-GivenName parameter), last name (-Surname parameter), and display name (-DisplayName parameter).

Although running this command would create the user account, there would be two caveats. First, the account would be disabled. Second, the account wouldn't have a password associated with it, which is required in most domains.

To avoid having to enable the account and add a password separately, you can modify the New-ADUser command I showed you. To have New-ADUser automatically enable the account, you need to

specify the *-Enabled \$true* parameter in the command. An enabled account requires a password, so you also need to specify the password in the command.

To provide a password, you can use the *-AccountPassword* parameter. However, you can't simply enter the password in plaintext on the command line. This parameter requires that the password be passed in as a secure string (i.e., have a data type of *SecureString*). There are two ways to convert the password into a secure string, both of which involve using a variable.

The first method uses the *ConvertTo-SecureString* cmdlet, which converts plaintext strings to secure strings. For example, if I want to convert the password *P@ssw0rd12* into a secure string and assign it to the *\$pwd* variable, I'd run the command:

```
$pwd = ConvertTo-SecureString -string "P@ssw0rd12" `
    -AsPlainText -force
```

This isn't the safest method for providing a password, because someone could be looking over my shoulder as I type this command. A safer way is to have the *New-ADUser* command prompt me for the password and mask the password as I type it. This can be done with the *Read-Host* cmdlet and its *-AsSecureString* parameter:

```
$pwd = Read-Host -AsSecureString
```

After this command runs, I'll see the familiar *\** character as I type my password. After typing the password, I'll need to press Enter.

Now that the password is stored as a secure string in the *\$pwd* variable, I can pass it to the *New-ADUser* cmdlet as follows:

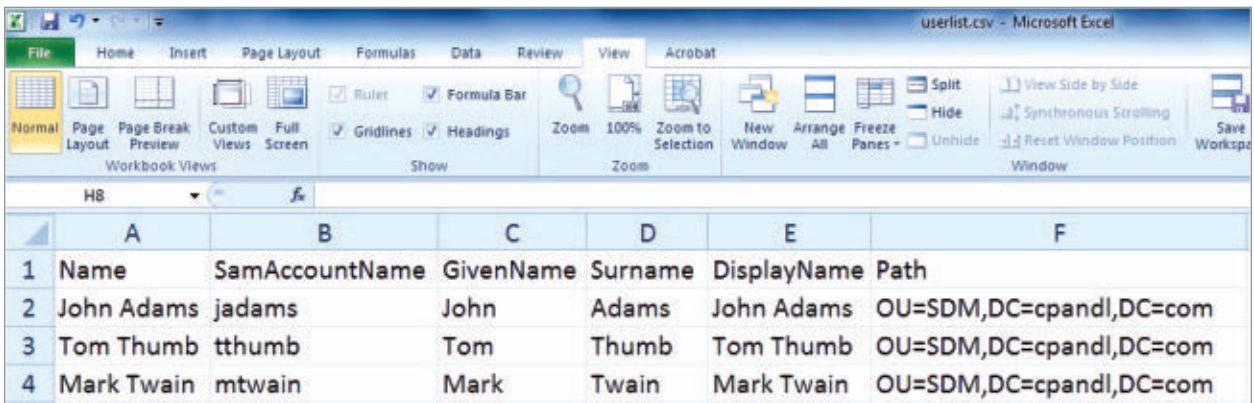
```
New-ADUser -Name "Bill Smith" -SamAccountName "bsmith" `
    -GivenName "Bill" -Surname "Smith" `
    -DisplayName "Bill Smith" `
```



```
-Path "OU=SDM,DC=cpandl,DC=com" `
-Enabled $true -AccountPassword $pwd
```

The command includes the `-Enabled` and `-AccountPassword` parameters to enable the account and securely associate a password with it.

Creating one user at a time is neat, but you might need to provision several users at the same time. This is where PowerShell can really shine. For example, suppose I need to create three user accounts. I can create a comma-separated value (CSV) file that contains the account information, then use the `Import-CSV` cmdlet to feed that information to the `New-ADUser` cmdlet. Figure 4 shows my CSV file, which is named `userlist.csv`.



	A	B	C	D	E	F
1	Name	SamAccountName	GivenName	Surname	DisplayName	Path
2	John Adams	jadams	John	Adams	John Adams	OU=SDM,DC=cpandl,DC=com
3	Tom Thumb	tthumb	Tom	Thumb	Tom Thumb	OU=SDM,DC=cpandl,DC=com
4	Mark Twain	mtwain	Mark	Twain	Mark Twain	OU=SDM,DC=cpandl,DC=com

In this file, notice that the column headers correspond to the parameter names provided in the previous `New-ADUser` command. Once again, this is intentional. When the CSV data is fed into the `New-ADUser` cmdlet, the cmdlet will pick up these parameter names from the PowerShell pipeline so that I don't have to specify them in the command. So, here's the command that I'd use to create the three user accounts:

```
Import-CSV -Path C:\data\userlist.csv |
New-ADUser -Enabled $true -AccountPassword $pwd
```

**Figure 4**

Using a CSV File to  
Create Several Users at  
the Same Time

As you can see, I piped the strings outputted by the Import-CSV cmdlet to the New-ADUser cmdlet. Because the pipeline understands that the column headers in the CSV file are parameter names and the rest of the rows contain the values, I only need to provide the -Enabled and -AccountPassword parameters. This is a great feature of the pipeline. It makes using PowerShell for these kinds of automation tasks that much more powerful.

## Managing Group Membership

Adding users or computers to groups is a common task in AD management. The AD module makes it relatively easy to perform this task. With the Add-ADGroupMember cmdlet, you can add one or more accounts to a group. For example, suppose I want to add the three users I just created to the Marketing Users group. The simplest way to do that is to run the command:

```
Add-ADGroupMember -Identity "Marketing Users" `
-Members jadams,tthumb,mtwain
```

In this command, I'm using the -Identity parameter to provide the name of the group. I'm also using the -Members parameter to provide the users' SAM account names. If you have multiple SAM account names, they need to be provided in a comma-separated list.

You can combine the operation to create the three users and the operation to add them to the Marketing Users group in a single command so that everything is done in one shot. However, the Add-ADGroupMember cmdlet doesn't support passing group member names into the pipeline. Therefore, you need to use the Add-ADPrincipalGroupMembership cmdlet if you want to leverage the pipeline. This cmdlet can take user, computer, or group objects as input from the pipeline and add those objects to the specified group.

Here's how to combine the user creation operation with the operation to add the new users to the Marketing Users group in one command:



```
Import-CSV -Path C:\data\userlist.csv |  
    New-ADUser -Enabled $true -AccountPassword $pass `  
    -PassThru | Add-ADPrincipalGroupMembership `  
    -MemberOf "Marketing Users"
```

Notice that I'm adding the `-PassThru` parameter to the `New-ADUser` portion of the command. This parameter tells `New-ADUser` to pass the user objects it creates to the pipeline. If this parameter isn't included, the `Add-ADPrincipalGroupMembership` cmdlet will fail.

Also notice that I'm using only the `-MemberOf` parameter to specify the group name in the `Add-ADPrincipalGroupMembership` portion of the command. The pipeline takes care of the rest, adding each of my three new users to the Marketing Users group.

So, using a single PowerShell command, I created three new users, put those users in my OU, gave them passwords, and added them to the Marketing Users group. Now let's look at some other common AD maintenance tasks that you can automate using PowerShell and the AD module.

## Resetting User Account Passwords

Occasionally users might need to have their user account passwords reset. You can automate this task with the `Set-ADAccountPassword` cmdlet, which is pretty straightforward. You use it to either change or reset an account password. If you change a password, you need to know the old password and supply the new one. If you want to reset the password, you only need to supply the new password. However, you need the Reset Password permission on the user object in AD to be able to perform the password reset.

Like the `-AccountPassword` parameter of the `New-ADUser` cmdlet, the `Set-ADAccountPassword` cmdlet leverages the `SecureString` data type for passwords, so you need to use one of the techniques I showed you to convert the plaintext passwords into secure strings. For example, suppose that I need to reset the password for the Tom Thumb user

account I created. After I store the new password as a secure string in the \$pass variable, I can run the command:

```
Set-ADAccountPassword -Identity "tthumb" `
-NewPassword $pass -Reset
```

In this command, I'm using the -Identity parameter to provide the SAM account name for the Tom Thumb user account. I'm also using the -NewPassword parameter with the \$pass variable to provide the new password. Finally, I'm specifying the -Reset parameter to tell the cmdlet that this operation is a password reset rather than a password change.

There's one additional task I want to perform: Toggle the flag on the Tom Thumb user account to force Tom to change his password at next logon. This is a common practice when you have to reset a user's password. I can accomplish this task by using the Set-ADUser cmdlet with the -ChangePasswordAtLogon parameter set to \$true:

```
Set-ADUser -Identity tthumb -ChangePasswordAtLogon $true
```

You might be wondering why I didn't use the pipeline to send the output of the Set-ADAccountPassword cmdlet to the Set-ADUser cmdlet to accomplish both operations in a single PowerShell command. Interestingly, I tried that and it didn't work. I presume there's some limitation in the Set-ADAccountPassword cmdlet that prevents the single command from succeeding. In any case, it's easy enough to toggle the flag using the Set-ADUser command I just showed you.

## Searching for Active Directory Objects

Another common AD task is to search for AD objects that meet certain criteria. For example, you might need to find all the computers running a certain Windows OS in an AD domain. The Get-ADObject cmdlet is the best cmdlet to use for LDAP searches. For example, if I

want to find the computers running Server 2008 R2 in my cpandl.com domain, I'd use the command:

```
Get-ADObject -LDAPFilter `
    "(&(operatingSystem=Windows Server 2008 R2 Enterprise)`
    (objectClass=computer))" -SearchBase "dc=cpandl,dc=com" `
    -SearchScope Subtree
```

This command uses three parameters to get the job done: `-LDAPFilter`, `-SearchBase`, and `-SearchScope`. The `-LDAPFilter` parameter takes a standard LDAP query as input. In this example, I'm querying for all computer objects that have their `OperatingSystem` attribute set to *Windows Server 2008 R2 Enterprise*. The `-SearchBase` parameter tells the cmdlet where to start the search in the AD hierarchy. In this case, I'm searching from the root of the cpandl.com domain, but you can easily limit the search to a particular OU if desired. The `-SearchScope` parameter tells the cmdlet whether to recurse all containers underneath the search base to find the specified objects. In this case, I'm using the `Subtree` option so that the cmdlet searches all the containers underneath it.

When I run this command, it'll display the objects that meet my criteria. Alternatively, I could pipe the results to other cmdlets to do something with those objects.

Note that for large searches, you might find it useful to use the `-ResultPageSize` parameter to control the paging of search results. I typically set this parameter to 1000 to tell the `Get-ADObject` cmdlet to return 1,000 objects at a time. Otherwise, you might find that you don't get all the results you expected because the number of objects being returned exceeds the maximum policy set for a single search request.

Another cmdlet that Microsoft provides related to searching is `Search-ADAccount`. This cmdlet is especially useful for searching for a variety of preset conditions, such as disabled accounts, accounts

with expired passwords, and accounts that are locked out. For example, the following command finds all user accounts that have expired passwords in my SDM OU:

```
Search-ADAccount -PasswordExpired -UsersOnly `
  -SearchBase "OU=sdm,dc=cpandl,dc=com" `
  -SearchScope OneLevel
```

In this command, I'm using the `-PasswordExpired` parameter to indicate that I'm looking for accounts with expired passwords. The `-UsersOnly` parameter tells the cmdlet to search user objects only (i.e., exclude computer objects). As I did in the previous example, I'm using the `-SearchBase` and `-SearchScope` parameters to tell the cmdlet where to search. However, in this case, I'm using the `OneLevel` option to search within the immediate OU only (i.e., not within any child OUs) to find accounts with expired passwords.

## Only Scratched the Surface

I've only scratched the surface of the capabilities of the AD module, but I hope you have a sense of the power contained therein. As I mentioned previously, there are more than 70 cmdlets in this module. Areas that I didn't touch on include deleting objects using the *Remove-*cmdlets, restoring deleted objects using the *Restore-ADObject* cmdlet, and modifying User Account Control (UAC) properties on user objects with the *Set-ADAccountControl* cmdlet. If there's an AD administration task you need to perform, there's likely a cmdlet that can handle the job. ■

# Data Protection Manager for Virtualized Workload Protection

## Take advantage of this under-utilized component of System Center 2012 SP1

**M**icrosoft System Center 2012 Data Protection Manager (DPM) is the most unused component of System Center. But for most organizations, this component can provide huge benefits. This is true for organizations that use Microsoft services such as Exchange, SharePoint, and SQL Server—and especially for those that use Hyper-V as their virtualization technology (an ever-growing percentage of organizations since the release of [Windows Server 2012](#)). In this article, I'll cover DPM's capabilities as they relate to Hyper-V. I'll also discuss considerations that relate to using Server Message Block (SMB) for virtual machine (VM) storage and clusters, using cluster shared volumes (CSVs), and live-migrating VMs while also using the best design to protect your workloads.

### DPM Hyper-V Protection 101

Before I discuss DPM and Hyper-V, I want to review Hyper-V's native backup capability. Windows uses Volume Shadow Copy Service (VSS) to ensure that backups are in an application-consistent state. Therefore, protected data will be usable when a data restore is needed. Without VSS, a backup process that has just backed up a running data file (e.g., a SQL Server database) would first need to back up a locked file. But then there would be no way of knowing that the data that was written to disk was in a consistent state. SQL Server might



### John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Mastering Hyper-V 2012 R2 with System Center and Azure* (Wiley).



**Email**



**Twitter**



**Website**



**Blog**

have been halfway through writing data, so the file might be corrupt and unusable in a restore process. VSS solves this issue by actively involving the application during a backup process.

Nearly all enterprise applications provide a VSS writer. This VSS writer allows application developers to define the actions that are required to ready an application's disk-based data for backup, make sure that all of the data in memory is flushed out to disk, and then pause any future writes to disk until the data snapshot is complete. All VSS writers that have been registered with an OS are called during a VSS backup to ensure that all of the data on disk is in an application-consistent state.

Why does this matter with regard to Hyper-V? Consider this: If a VM is backed up from the Hyper-V host, the virtual hard disk (VHD, or newer VHDX) files are backed up. If a snapshot of the current state of a VHD or VHDX was simply taken and backed up, then the current state of that VHD or VHDX is unlikely to be in a consistent state. The OS that is running inside the VM has no way of knowing that a backup is being performed. The only way to ensure a good backup is to perform an offline backup of the VM. Fortunately, Hyper-V addresses this situation.

In an enlightened VM, the OS that's running in the VM knows that it's virtualized and can communicate with the virtualization host. This communication occurs primarily through the Hyper-V integration services, of which there are several. One is the Backup (Volume Snapshot) integration service, which is the key to application-integral backups of Hyper-V VMs.

When a backup of a Hyper-V VM is taken, the VSS request is passed, via the Backup integration service, to the OS that's running within the VM. The passing of the VSS request calls all VSS writers that are registered within the VM's OS. The data of the VM on the VHD or VHDX is readied for backup, so when that VHD or VHDX is backed up at the host level via the Hyper-V writer, the content is in an application-consistent state.

DPM fully utilizes this VSS pass-through Hyper-V capability for its online protection of Hyper-V VMs. However, there are requirements for an online backup of VMs:

- The Backup integration service must be enabled, which means that the OS running in the VM must support Hyper-V integration services.
- The Windows guest OS must support VSS (Windows 2003 or later).
- Dynamic disks must not be present within the VM.
- All volumes must be NTFS—even when Microsoft Application Virtualization (App-V), which might create a non-NTFS volume, is used.
- The VM must be running.
- VSS storage assignment for the volumes must not be modified.
- If the VM is part of a cluster configuration, then the cluster resource group must be online.

If any of these requirements are not met, such as with a Linux VM or Windows 2000 system, then an offline backup is performed. This backup places the VM in a saved state while the snapshot is taken, then resumes the VM. This causes a period of unavailability for the VM during the backup. However, this period should be only about 30 seconds in most environments.

When creating a new DPM protection group or adding a VM to an existing protection group, the type of backup that will be used is shown, and the basic process is as follows:

1. Launch the DPM Administrator console and select the Protection workspace.
2. Choose the New action.
3. Click Next to open the Introduction wizard.
4. Choose Servers as the type of protection group, and then click Next.
5. Choose the Hyper-V server that you want to protect and then expand the HyperV navigation node to see all the VMs that can

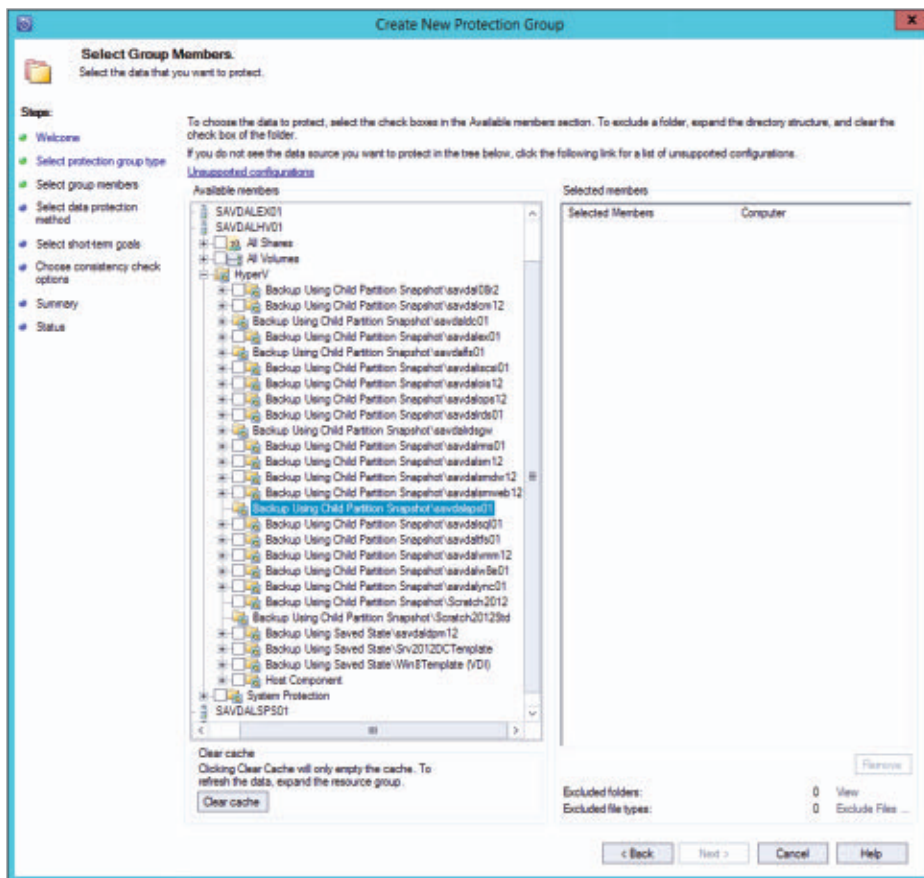


be protected, as Figure 1 shows. Note that there are two options for the backup:

- Backup Using Child Partition Snapshot—an online backup (for VMs that meet the previously listed requirements)
- Backup Using Saved State—an offline backup

**Figure 1**

Selecting Protected VMs as Part of the Protection Group



6. Choose the VMs that you want to back up, and then click Next.
7. Choose a name for the protection group. For System Center 2012 Service Pack 1 (SP1) DPM, you can also save VM backup data to online storage, such as Windows Azure storage.
8. Complete the remaining steps to configure retention time, times to create recovery points, how to create the initial

replica (i.e., over the network or via removable media), and storage allocation.

When Hyper-V protection is in place, you can restore either the entire VM or specific files from the VM. This feature is useful because mounting VHDs is native to Windows Server 2008 R2 and later. Use the Recovery workspace, choose the VM, choose a date and time, and double-click the VM VHD in the details area. Continue to drill down into the volumes, folders, and files from the VM and restore the items that you need.

## Special Hyper-V Protection Considerations

There are some special considerations when using DPM to protect Hyper-V VMs, especially when using Windows Server 2012 Hyper-V. Think about protecting a typical server OS: The agent runs within the OS, and the DPM server can always communicate with that agent to protect data such as files for the SQL Server database. The same is true for a VM running on Hyper-V. The DPM server can communicate with the DPM agent on the Hyper-V host to retrieve data for the protected VM.

What about live migration of a VM within a cluster? The VM can be on Server A one day and on Server B the next, so how can DPM know how to continue protecting the VM? For VMs in a Windows Server 2012 cluster, continued protection—even when VMs are migrated live between nodes—is supported, providing that the following steps are performed prior to protecting the VMs:

1. The DPM agent has been installed on all nodes in the cluster.
2. The System Center 2012 Virtual Machine Manager (VMM) console has been installed on the DPM server. (Make sure that the VMM console is the same version as the VMM server that is managing the Hyper-V cluster.)
3. The following Windows PowerShell command has been run from an elevated DPM PowerShell prompt:

```
runSet-DPMGlobalProperty -DPMServerName <Name of the DPM
server> -KnownVMMServers <Name of the VMM server>
```

#### 4. The DPMVMMHelper server has been started.

Speaking of clusters: If CSVs are used in a large cluster (i.e., 8 nodes or more) with many VMs (e.g., 400 VMs), a scan of the environment could take hours to complete when you want to configure protection. However, with Windows Server 2012 and Service Center 2012 SP1 DPM, you no longer need to enable serialization. However, it's still better to use a hardware VSS provider when possible, to avoid a performance impact on protected VMs.

Another new feature of Windows Server 2012 is support for storing VMs on SMB 3.0 file shares. This capability adds additional complexity to DPM protection. Fortunately, as part of making SMB an enterprise-quality protocol, Microsoft provides VSS support for SMB file shares. DPM uses this support when protecting VMs that are stored on SMB 3.0 file shares. To ensure proper protection of such VMs, the DPM agent must be installed on the SMB file server (or servers, if the SMB share is in a cluster) and the File Server VSS Agent Service (a role service of Windows Server 2012) must be enabled.

Windows Server 2012 also supports shared-nothing live migration, which allows VMs to be moved between Hyper-V hosts that aren't part of a cluster (or that are part of different clusters). As long as both the source and target Hyper-V hosts are managed by the same VMM server, DPM protection will continue. (Note the common thread: VMM is crucial for continued protection when VMs move between hosts.) When dealing with shared-nothing live migration, two scenarios can affect what happens to VM protection:

- The VM is stored on an SMB file share, so the only movement is the VM memory and state (not the storage). There is no real change to what DPM is protecting, and the protection continues without interruption.

- The VM is stored on non-common storage, which means that it requires a storage migration as part of the shared-nothing live migration process. Although DPM can still recognize that the VM has moved, it requires a consistency check (i.e., a block-by-block comparison of the source protected data and the DPM replica, to look for possible inconsistencies) on the storage if protection is to continue. This requirement results in some overhead on the system.

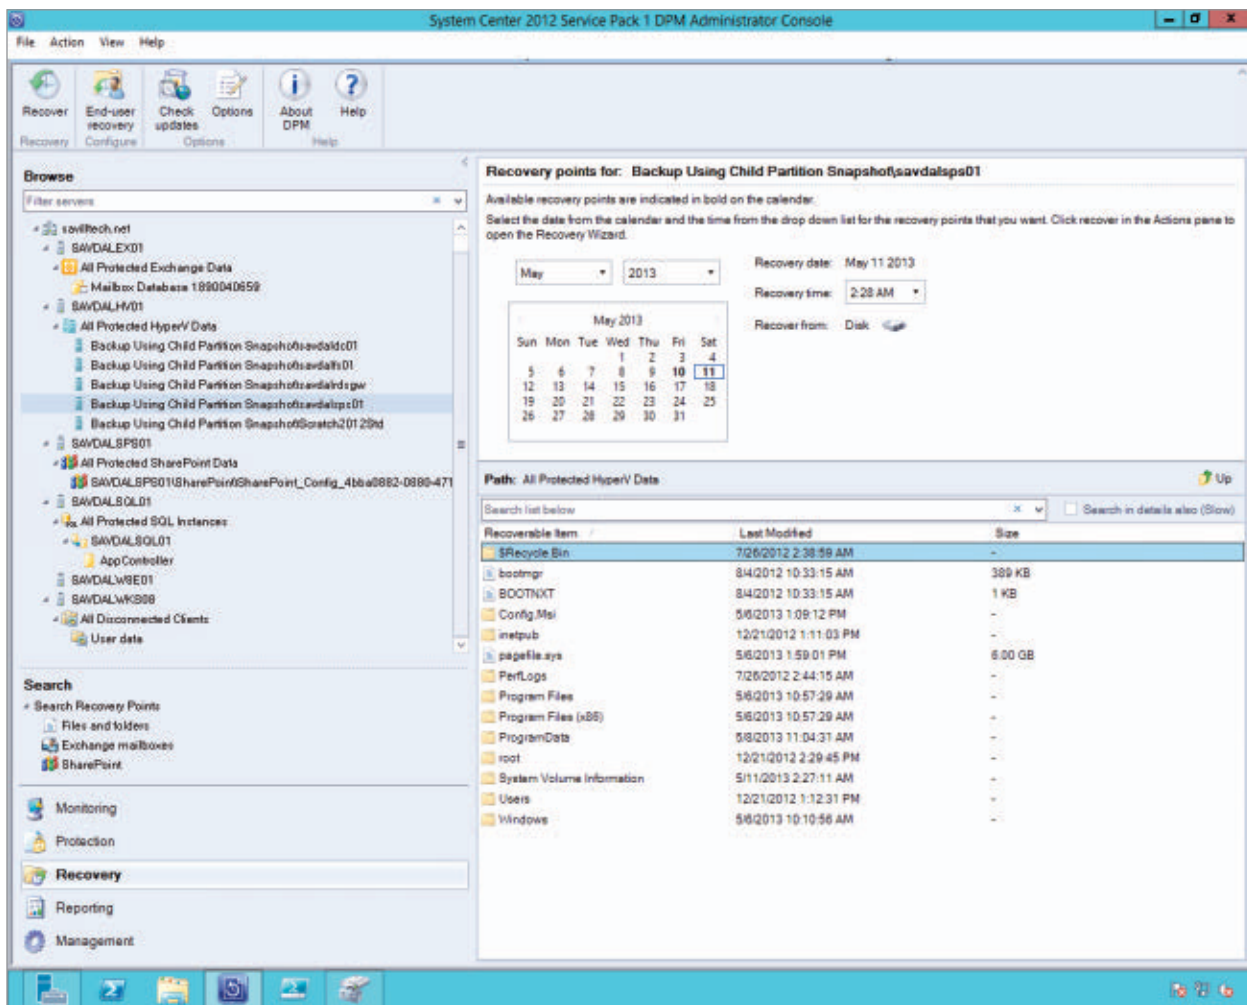
Although some extra steps are required for certain advanced scenarios, you can take comfort in the fact that all scenarios are possible. In any given situation, DPM can fully protect your Hyper-V environment.

## How to Protect Hyper-V VMs

How can you protect your Hyper-V VMs? You might think that simply choosing the VMs as part of a protection group at the Hyper-V host level would be the simplest method. However, this might not be the best option.

Consider the fact that DPM has application-level protection of workloads such as SQL Server, Exchange, and SharePoint, enabling application-level recoveries of databases, mailboxes, or SharePoint items. Should you protect a SharePoint server that is running inside a Hyper-V VM by protecting the VM at the Hyper-V host level? If you need to restore data, you can restore the entire VM or files from the file system, but none of that restore is in SharePoint terms, and directly restoring SharePoint items isn't possible. But if that same SharePoint server is protected by installing the DPM agent within the VM and protecting SharePoint from within the VM as part of a protection group, then you can restore individual items from SharePoint.

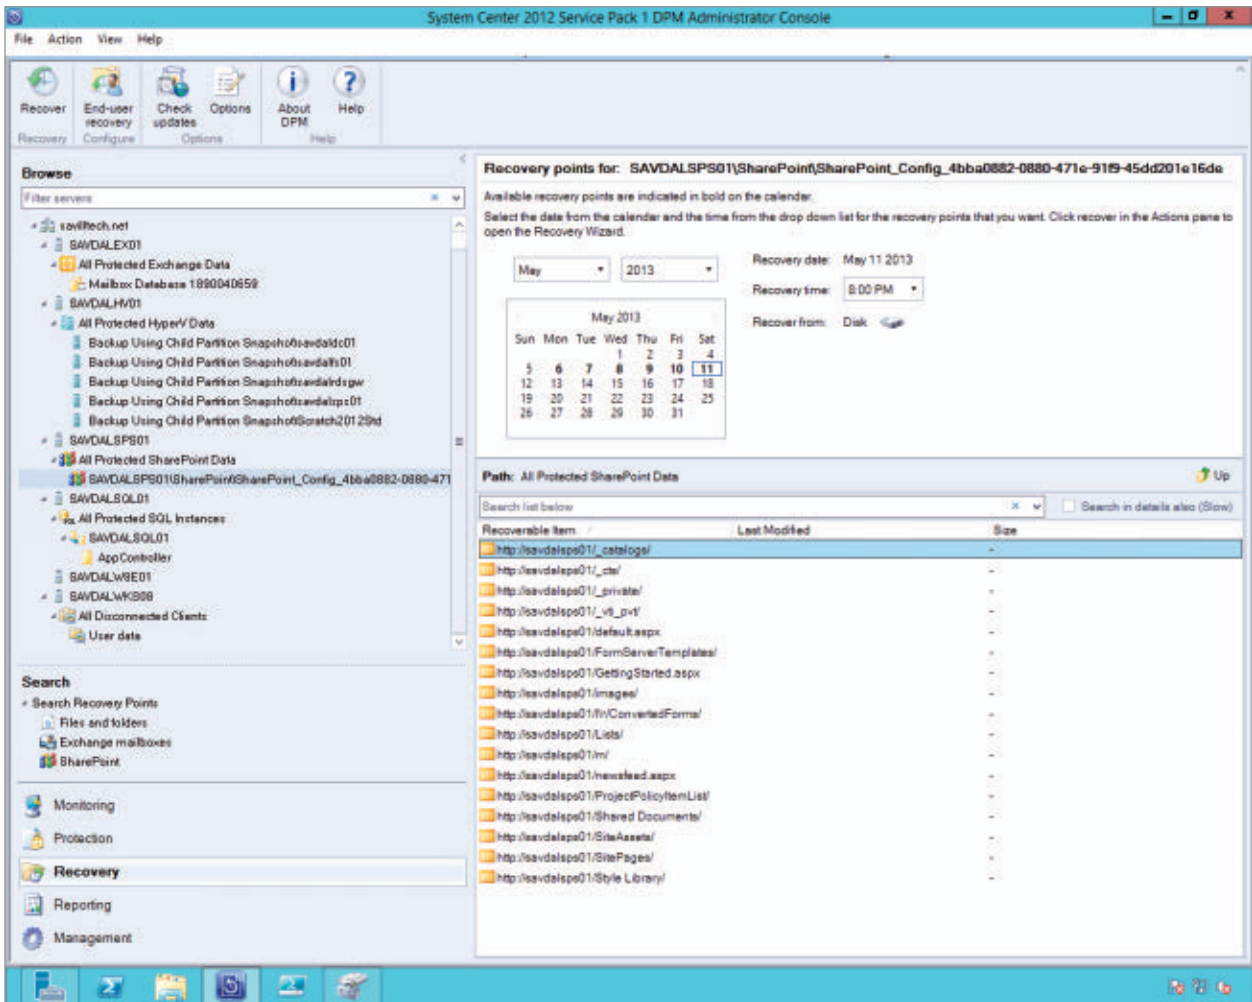
Figure 2 and Figure 3 show the same VM. Figure 2 shows the restore capability when the VM is protected at the Hyper-V host level;



**Figure 2**  
Hyper-V Host-Level  
Restore Granularity

Figure 3 shows the restore capability when SharePoint is protected within the VM. As you can see in these figures, there's a big difference between the two options.

Backing up VMs at the Hyper-V host level is not always the best possible option. The decision regarding how to back up your VMs should be based on how you want to perform restoration of your data. If you want application-aware restoration, then you need to run the DPM agent within the VM and use that agent to protect the



**Figure 3**  
Hyper-V VM  
Application-Level  
Restore Granularity

workloads within the VM. Otherwise, backing up at the Hyper-V host level is sufficient.

## Not Just for Hyper-V

I strongly urge you to consider DPM—not just for Hyper-V, but also for other Microsoft workloads. Don't take the path of least resistance, which would be to simply protect all VMs at the Hyper-V host level. If possible, use more granular, application-aware protection



## Video



John Savill provides an overview of DPM with Hyper-V



by installing the DPM agent within your VMs, and use protection at the Hyper-V host as a second choice. For more information, see the accompanying video.



# Office 365 Message Encryption

## Protect your email against the spooks

**T**he [PRISM revelations](#) have had a huge effect on the general public and on the cloud community in particular. No one likes the idea of their private communications being intercepted and analyzed, and the fact that three-letter agencies were busily tapping into various links to capture data came as an unpleasant and unwelcome surprise—especially to companies considering cloud deployments.

As a long-time consumer of cloud services, I don't have a problem if governments want to read my email—but then again, my data is probably less interesting and less worth intercepting than data from large companies, international bodies, and the like. This might explain the efforts of cloud service providers to reassure customers and potential clients that it's safe to hand over control of data—which is where we come to the [new Office 365 Message Encryption feature](#), a replacement of the older [Exchange Hosted Encryption \(EHE\)](#) service, announced by Microsoft on November 21, 2013. The new service will be available for purchase in the first quarter of 2014, and customers who use EHE will be upgraded to use Office 365 Message Encryption at around the same time.

Office 365 Message Encryption is included in the [Windows Azure Rights Management suite](#) (RMS), so it's free if you subscribe to Office 365 plans E3 or E4. Rights Management is available as an add-on for other Office 365 plans and costs \$2 per user per month (in the United States). Azure RMS also works in a hybrid Exchange environment if messages from the on-premises mailboxes pass through Exchange Online Protection (EOP).

### How Office 365 Message Encryption Works

The new Message Encryption feature works by intercepting messages as they pass through the transport pipeline and encrypting all



### Tony Redmond

is a senior contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 Inside Out* (Microsoft Press) and *Microsoft Exchange Server 2013 Inside Out: Mailbox and High Availability* (Microsoft Press).



**Email**



**Twitter**



**Blog**

messages that meet whatever criteria the administrator sets. Users don't have the option to encrypt selected messages—the determination is made through transport rules, so you have to come up with a set of criteria to identify outbound messages that need to be protected. For example, all messages that have “Confidential Data” in their subject line, messages sent to a particular domain, or even messages sent to a particular user would be encrypted. Different transport rules can be created to handle different circumstances, all of which result in encrypted outbound messages.

Office 365 Message Encryption extends the “modify the message security” predicate to add the “Apply Office 365 Message Encryption” action; “Remove Office 365 Message Encryption” is also a supported action. Applying encryption is one of many actions that can be invoked by transport rules, so it can be mixed with all manner of other rule conditions and actions to perform whatever processing is required. Because encryption is incorporated into transport rules, it's guaranteed that no message will be dispatched before it is evaluated and, if it meets the criteria set for a rule that requires encryption, it's also guaranteed that each message is protected before dispatch.

Messages are 2048-bit encrypted using SHA-256 with the private key of the Office 365 tenant domain. Recipients have no knowledge of this key, so when they receive a message, they'll see that it contains an encrypted attachment and instructions for viewing the content. Because all the email system has to do is transport and deliver the encrypted attachment, you can expect the solution to work for messages sent to all other major commercial and consumer email systems.

Clicking the attachment opens a browser window connected to a page on the Office 365 Message Encryption portal. Companies can customize this page with a logo and some directive text to tell end users what to do. The user will then authenticate using a Windows Live or Office 365 ID before the content is decrypted and presented in an Outlook Web App–like interface. Any reply sent back to the originator is automatically encrypted.

## Deploying Office 365 Message Encryption

Overall, I like Microsoft's approach in designing Office 365 Message Encryption. Microsoft leveraged some of its assets from Windows Azure RMS, transport rules, and OWA to assemble a solution that protects confidential messages against anyone who penetrates a mailbox or intercepts email en route. Making Office 365 Message Encryption an administrator-driven process means that users don't get to vote whether their messages are encrypted. It also means that you avoid the key issuance and maintenance that user-centric message encryption involves, something that often works great for small groups of users but rapidly becomes a royal pain when rolled out into general use.

Despite the positive vibes, I wouldn't rush to deploy Office 365 Message Encryption, at least not without doing some thinking first. Some work has to be done to figure out how to identify confidential messages to transport rules and then how to integrate encryption into the transport rule pipeline so that it doesn't interfere with any other rules processing such as those that implement data loss prevention. You then have to consider user education and figure out what needs to be done to advise external correspondents on how to handle the encrypted messages they will now receive from your company. It's also important to realize that this implementation, in terms of accessing the encrypted message content, is strictly browser-based for now and that there's no offline client support.

## More Knowledge Is Key

Like any new feature, an imperfect state of knowledge exists regarding operational issues, such as gaining access to important encrypted messages received by people who leave the company. It's also important to consider how the use of Office 365 Message Encryption affects your Plan B—your back-out plan in case cloud services don't work for your company. The answer here seems emphatic: If you give up your Office 365 tenant, you give up the Rights Management private key, and any messages encrypted using this key become inaccessible. ■

# FAQ

## Answers to Your Questions



John Savill

**Q:** Does the FTP server that's included with Microsoft IIS support FTP over SSL? If it does, how can I configure it? Can I configure it to secure only the exchange of the FTP user credentials?



Jan De Clercq

**A:** The IIS FTP server supports FTP over SSL, starting with the IIS 7.0 web server that's bundled with [Windows Server 2008](#). To enable FTP over SSL, you should first make sure that you have a valid SSL certificate configured for your web server. You can create a self-signed certificate or obtain a certificate from your enterprise Certification Authority (CA) or from a commercial CA. You can configure an SSL certificate using the Server Certificates option that shows up in the center pane of IIS Manager when you select your web server object.

You must also allow the use of SSL when you enable the FTP protocol for your website. To enable FTP from IIS Manager, select the website and click Add FTP Publishing in the Actions pane. Next, in the *Bindings and SSL Settings* section, make sure that you select Allow SSL in the SSL section. Finally, in the *Authentication and Authorization Information* dialog box, you typically select Basic in the Authentication section, select the Specified Users option in the *Allow access to* drop-down list, and enter the FTP user logon account in the accompanying text box in the Authorization section.

You can then further configure FTP over SSL using IIS Manager's FTP SSL Settings feature, which you can find in the center pane, both on the web-server level and website level. This is the feature you'd use to configure IIS to secure only the FTP credential exchanges

using SSL. It's important that you configure the same settings on both the web-server level and website level. If you don't do so, you'll get "conflicting TLS setting" error messages when you try to connect to your FTP site.

To change the FTP SSL settings for your website, navigate to your website from IIS Manager and double-click FTP SSL Settings in the center pane to open the FTP SSL Settings dialog box. From the SSL Certificate list, select the certificate that you want to use for securing your FTP connections. Under SSL Policy, you can select one of the following options.

***Allow SSL connections.*** Choose this option if you want your FTP server to support both non-SSL and SSL connections.

***Require SSL connections.*** Select this option if you want to enforce the use of SSL encryption for all FTP communications.

***Custom.*** Choose this option if you want to configure a different SSL policy for the FTP control and data channels. After you select it, click the Advanced button. In the Advanced SSL Policy dialog box, you can select the SSL policies. For the FTP control channel, your options are the following:

- *Allow.* Allows SSL for the control channel, meaning that SSL isn't required.
- *Require.* Enforces the use of SSL for the control channel.
- *Require only for credentials.* Requires that only the FTP user credentials have to be sent using SSL.

For the FTP data channel, your options are the following:

- *Allow.* Allows SSL for the data channel, meaning that SSL isn't required.
- *Require.* Enforces the use of SSL for the data channel.
- *Deny.* Denies the use of SSL for the data channel.

In your case, you'd select the Custom option. Then, in the Advanced SSL Policy dialog box, you'd choose *Require only for credentials* for

the FTP control channel and *Allow* for the FTP data channel. Remember that after you configure these settings for your website, you must also configure the same settings for your web server.

—Jan De Clercq

**Q:** Is it true that Windows Server 2012 lets you securely import PKCS#12 content without needing to transfer the associated protection password to all administrators involved in the import operation?

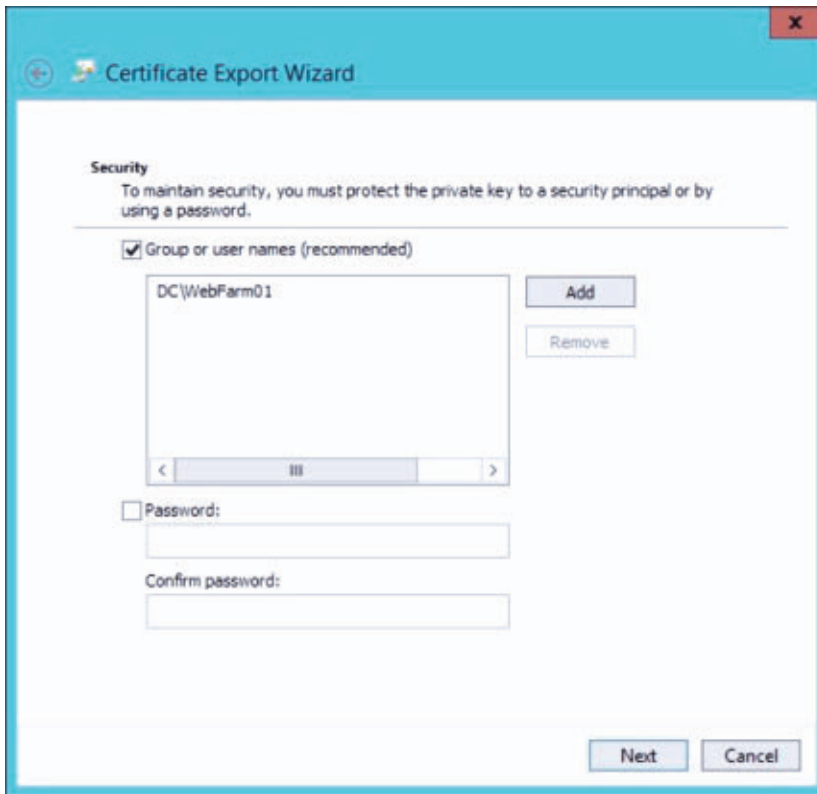
**A:** Yes. Windows Server 2012 and Windows 8 support a new feature that lets you use an Active Directory (AD) user or group account to protect the certificate and private key, both of which are contained in a PKCS#12-formatted file. This feature is useful for exporting and importing digital certificates and private keys because you don't need to share a password with all the parties involved.

To support this feature, the Certificate Export Wizard in Server 2012 and Windows 8 provides a new security option that lets you secure the .pfx file using an AD user or group account. The new *Group or user names (recommended)* option appears in the Security page of the wizard, as Figure 1 shows.

When you select this option, you can add an AD account using the Add button. The Certificate Import Wizard's logic has been changed to automatically detect when you're logged on with a valid AD account and unlock access to the certificate and private key in the .pfx file, without prompting you for a password.

Under the hood, this new feature still uses a password that's automatically generated to protect the .pfx file. This password is encrypted using the Windows Data Protection API (DPAPI) and added to the .pfx file.

When a protected .pfx file is imported, Windows checks whether the user or computer account that's trying to import the file is on the list of accounts that was configured during the creation of the .pfx file. If



**Figure 1**  
Security Page in the  
Certificate Export  
Wizard

that's the case, Windows automatically unprotects the password and gives access to the protected certificate and private key content.

This new feature only works when you export the certificate and private key from a Server 2012 or Windows 8 machine that's an AD domain member. In addition, the machine on which you import the certificate and private key must be joined to a domain where a Server 2012 domain controller (DC) is available.

You can invoke this new feature from the command line when you use either the [Windows PowerShell](#) `Export-PfxCertificate` cmdlet or the `Certutil` command-line utility to create a .pfx file. Both tools now support the `-ProtectTo` parameter, which lets you specify the AD account you want to use to protect the certificate and private key.

—Jan De Clercq



**Q:** How can I set the node that should lose its vote for the Windows Server 2012 R2 50/50 vote split?

**A:** Windows Server 2012 R2 introduces functionality to enable a partition of a cluster to keep running even when there's a 50/50 vote split. In the event of an even number of nodes with no witness configured (which, technically, should never happen in Server 2012 R2 if you always configure a witness), one node has its vote automatically removed to make the cluster have an odd number of votes again. This would allow one partition of the cluster to have a majority number of votes in the event of a split.

Instead of letting this node be randomly selected by the cluster, it's possible to specify the node that should have its vote removed, should there be an even number of node votes with no witness. To set it up so that the node is specified, set the `LowerQuorumPriorityNodeId` attribute of the cluster to the ID of the node that should lose its vote. For example, this is what I entered:

```
(Get-Cluster).LowerQuorumPriorityNodeId = (Get-ClusterNode savda1hv20).NodeInstanceId
```

—John Savill

**Q:** Why were snapshots renamed to checkpoints in Windows Server 2012 R2 Hyper-V?

**A:** There was always an inconsistency in naming. Whereas Hyper-V used the “snapshot” nomenclature, System Center Virtual Machine Manager used the “checkpoint” nomenclature. To provide consistency and avoid confusion, snapshots in Hyper-V have been renamed to checkpoints. ■

—John Savill

# Product News for IT Pros

## Vision Solutions Extends Database Replication Capabilities

Vision Solutions announced the release of Double-Take Share 5.2, the latest version of the company's cross-platform data-sharing solution. Double-Take Share improves organizational decision-making and productivity by making it easy to replicate mission-critical data across multiple databases, incompatible applications, and differing hardware and OS platforms. Version 5.2 adds full support for the latest Windows and SQL Server technologies, increases performance, and adds support for new data types and OS and database versions. Double-Take Share 5.2 uses a simple, point-and-click graphical UI to eliminate the need for scripting or programming to move data between databases, which improves database administrator productivity. For more information, visit the [Vision Solutions website](#).



## SolarWinds Enhances Security Offerings

SolarWinds recently announced enhancements to a few of its security management solutions, including SolarWinds Log & Event Manager, SolarWinds Firewall Security Manager, and SolarWinds Patch Manager. SolarWinds Log & Event Manager now provides additional automatic scheduling, distribution, and notification of security and network activity, giving IT pros greater assurance that regular reviews are being performed and ensuring that strong security monitoring and compliance management practices are met. The latest version of SolarWinds Firewall Security Manager, which automates heterogeneous firewall configuration and change management, provides integration with SolarWinds' centralized management console. Now, administrators can quickly identify high-risk firewalls, recent and upcoming configuration changes, and firewall compliance status in



SolarWinds' single pane of glass for managing network, systems, and security. To ensure that patch management processes are optimized, SolarWinds Patch Manager has been enhanced with step-by-step wizards, enabling users to be up and running quickly and significantly reducing the time it takes to get vulnerable machines patched and protected. Learn more at the [SolarWinds website](#).



### **Dell and Microsoft Strengthen Public Cloud Alliance**

Dell announced that it's building on its strategic alliance with Microsoft to deliver Windows Azure to Dell customers worldwide through the Dell Cloud Partner Program. This expands on the previously announced alliance providing Application Development Services on Windows Azure. With the evolving Dell/Microsoft relationship, current and future customers will have even more choice and flexibility when pursuing and planning public cloud infrastructures. Acting as a single-source supplier through the Dell Cloud Partner Program, Dell will offer customers a central point of solution integration, control, and direct support, lessening the complexity and challenges of deploying cloud environments. For more information, visit [Dell](#).



### **Symantec Enables Data Centers to Adopt SSDs**

Symantec announced a new version of its Storage Foundation software that lets data centers leverage SSDs in ways that allow customers to access mission-critical data and applications 400 percent faster than over traditional SANs. It is also the only offering to provide these benefits regardless of which storage hardware components are in place. Customers are therefore free to choose any storage infrastructure provider. Storage Foundation 6.1's functionalities and benefits include a vendor-agnostic intuitive caching layer (enabled by Symantec's SmartIO technology) that detects critical application workloads and caches only the hot data on local SSDs, and Flexible Storage Sharing (FSS) technology that enables servers to access remote data as if it were from local storage. This allows organizations to reduce storage

costs by up to 80 percent using commoditized storage hardware, while helping to ensure that all data is replicated, protected, and available. For more information, check out the [Symantec website](#).

## First Cloud-Based Facilities Management for Microsoft Office 365



SP Marketplace announced SP Facilities Management for Office 365 and SharePoint 2013, the first computerized maintenance management system (CMMS) that runs on Office 365 and SharePoint Online. Unlike traditional database applications, SP Facilities Management is fully customizable by business users (without coding) to fit specific organization needs. SP Facilities Management is a simple-to-use yet powerful management program for the maintenance of equipment and facilities. SP Facilities Management lets business managers track regular maintenance and emergency work, send work orders to the maintenance staff on their mobile phones, and track the progress of all outstanding and completed jobs. It also includes a service-request portal for employees or customers to track the status of open requests, access a knowledge base, and find relevant how-to documents. Leveraging Office 365 and SharePoint, SP Facilities Management integrates documents, email correspondence, calendars, and task tracking with work orders, facilities, and equipment. Learn more at [SP Marketplace](#).

## Netwrix Brings Free Password Management to Small Businesses



Netwrix announced the availability of Netwrix Password Manager as a completely free offering to businesses with up to 100 users. Netwrix Password Manager provides simplified password management that's self-serviced so that users can reset passwords and unlock their own accounts after successfully answering identity verification questions. The product also provides logon and web-based management, meaning that users can reset their passwords directly from a Windows logon screen or through a web browser. This allows organizations to

implement strong password policies in Active Directory (AD) to meet regulatory compliance requirements and address identity management challenges, without the risk of increasing the workload of Help desk personnel. For more information, visit the [Netwrix website](#).



### **LenovoEMC Announces Four-Bay Desktop NAS**

LenovoEMC recently announced the worldwide availability of the new LenovoEMC px4-400d Network Storage, a high-performance four-bay desktop NAS that delivers advanced data protection as well as local and remote content sharing, utilizing the simplicity and power of the LenovoEMC LifeLine OS. The LenovoEMC px4-400d uses hot-swappable 7200rpm SATA-III Server Class drives in a compact form factor, with up to 16TB in storage capacity (4 × 4TB) and up to 25 percent higher data read and write performance than the previous generation of product. Compatible with Windows, Mac, and Linux devices, the new LenovoEMC px4-400d is powered by the LifeLine OS, which incorporates standard network storage capabilities and many advanced business features, including built-in cloud technology, versatile remote access, virtualization, data replication, device-to-device copy jobs, advanced RAID support, and SSD support. New hardware features include HDMI, eSATA, and USB 3.0 ports for versatility of use. For more information, visit the [LenovoEMC website](#). ■

Search our network of sites dedicated to hands-on technical information for IT professionals.  
[www.windowsitpro.com](http://www.windowsitpro.com)

**Support**  
Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.  
[forums.windowsitpro.com](http://forums.windowsitpro.com)

**News**  
Check out the current news and information about Microsoft Windows technologies.  
[www.winsupersite.com](http://www.winsupersite.com)

- EMAIL NEWSLETTERS**  
Get free news, commentary, and tips delivered automatically to your desktop.
- Cloud & Virtualization UPDATE
  - Dev Pro UPDATE
  - Exchange & Outlook UPDATE
  - Security UPDATE
  - SharePoint Pro UPDATE
  - SQL Server Pro UPDATE
  - Windows IT Pro UPDATE
  - WinInfo Daily UPDATE

**RELATED PRODUCTS**  
*Windows IT Pro* VIP  
Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.  
[windowsitpro.com/vip-premium-membership](http://windowsitpro.com/vip-premium-membership)

*SQL Server Pro*  
Explore the hottest new features of SQL Server, and discover practical tips and tools.  
[www.sqlmag.com](http://www.sqlmag.com)

*Dev Pro*  
Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at [DevProConnections.com](http://DevProConnections.com), where IT pros creatively and proactively drive business value through technology.  
[www.devproconnections.com](http://www.devproconnections.com)

*SharePoint Pro*  
Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.  
[www.sharepointpromag.com](http://www.sharepointpromag.com)

Advertiser Directory

**ManageEngine** ..... 2

**Netwrix** ..... 1

**RSA Conference** ..... 6

**StorageCraft** ..... 7

**Windows IT Pro**..... 17, 49, 50

Vendor Directory

**Apple** ..... 8, 9

**Dell**..... 84

**Dropbox** ..... 5

**Google** ..... 12, 13

**LenovoEMC** ..... 86

**Netwrix** ..... 85, 86

**SolarWinds** ..... 83, 84

**SP Marketplace** ..... 85

**Stardock**..... 21

**Symantec** ..... 84, 85

**Vision Solutions** ..... 83

